

## $\infty$ 10 0 $\cdots$ 0 0 $\square$ AD



COMPUTER SCIENCE TECHNICAL REPORT SERIES


UNIVERSITY OF MARYLAND COLLEGE PARK, MARYLAND 20742


The research described in this report was funded by the office of Naval Research under contract number N0b014-76C-0477.

This report is distinct in content from but similar in topic
to U.O.M. TR-495, entitled "The Causal Representation and Simulation of Physical Mechánisms", Nov. 1976.


The Declarative Representation and Procedural Simulation of Causality in Physical Mechanisms

Chuck Rieger and Milt Grinberg

> Computer Science Department
> University of Maryland
> College Park, Maryland 20742


#### Abstract

A theory of cause-effect representation is used to describe man-made mechanisms and natural laws. The representation, consisting of 10 link types that interconnect events into large declarative patterns, is illustrated on a relatively sophisticated device, the home gas forced air furnace. Next, a procedure and framework for translating the declarative description of a mechanism into a population of associatively triggerable computation units is described. The associative, or procedural, form can then be used to perform a discrete cause-effect simulation of the device. The delcarative to procedural translation, including a simulation trace, is shown for the furnace. Topics Of mechanism abstraction and mechanism invention are discussed, and the entire Mechanisms Laboratory is placed in the larger perspective of out cresearch into human problem solving. $$
R
$$


Keywords: cause-effect representation, declarative-procedural correspondence, simulation of physical systems, spontaneous computation, mechanisms invention

## 1. Introduction

Man-rade physical mechanisms provide an interesting domain in which to study human problem solving and cause-effect knowledge representation. In this paper, we describe and illustrate a representation framework which permits us to express and simulate the commonsense internal cause-effect structure of a variety of manmade mechanisms (both mechanical and electronic), as well as a variety of natural mechanisms (laws of physics, physiological mechanisms, etc).

Since our strategy has been to use the mechanisms domain as a medium for investigating human problem solving and cause-effect knowledge representation, we have developed our "Mechanisms Lab" within the existing framework of our Commonsense Algorithm (CSA) Project, a broader investigation of cognitive mechanisms. Following this strategy has lead us to mechanisms theory in which the "ground" representation of a mechanism is a declarative, cause-effect graph, but in which simulation is accomplished by transforming such a representation into procedural form, then executing it. Thus, in addition to the theory of representation and simulation, we feel the technique provides an interesting case study in declarative-procedural correspondence.

The discussion of our Mechanisms Laboratory is divided into five main parts: (1) background and theoretical framework of the mechanisms research, (2) cause-effect representation, (3) background and strategy for the simulation aspects of the project, (4) a simulation example, and (5) philosophy of the approach.

## 2. Motivation

Our motive in this mechanisms research is to understand better human problem solving and cause-effect knowledge representation. Man-made mechanisms provide an excellent medium for carrying out such research because mechanisms are in a sense final snapshots of the human problem solving process. By developing representations for the cause-effect notions inherent to all mechanisms as humans perceive them, we stand a good chance of also gaining insight into how humans encode the various princinias and physical laws that are evoked during the mechanism's invention or design.

Besides this theoretical relevance of studying mechanisms, there are some interesting practical applications of a cognitive theory of mechanisms representation and simulation. For example, a mechanisms theory can provide the basis of a potentially powerful CAI facility in which a mechanisms-naive user could interactively explore the design principles and behavior of a device. Provided the system is conversant in the same terms as the user (i.e. higher-level, symbolic terms, rather than lower-level numerical or parametric terms), such interaction can communicate concepts rather than details.

Many other specific applications exist. In one, the theory could provide the basis for medical applications, through the modeling of, e.g., human physiological mechanisms. In a second, the theory could be applied to provide a self-model in systems with a self-maintenance capabilities (martian rovers, 2001 vacuum cleaners?). In a third, the theory could provide the basis of interactive design systems in which engineers specify high level goals to the system, which then does the design. Completely automatic "mechanisms invention" is within the realm of possibility and, we feel, would closely resemble the behavior of current-day problem solving and program synthesis theories. Finally, a future application might be to use the representation as a sort of Dewey Decimal System of cause-effect pattern classification.

## 3. Background and Related Work

There are many ways to approach the description of a mechanism. Most in the past (e.g. [JW1], [LBR1], [RJSI], [SL1], and [WBLI]) have tended to be more analytical in their approach. In analytical simulations, the mechanism is typically described by parameterizing it in form suitable for a numerical simulator. The problem with this approach is that the representation of a mechanism is very different from the human cause and effect knowledge of the mechanism.

We are by no means the only ones addressing the issues of symbolic modeling or simulation of physical systems. Two other notable examples are the MYCIN project and an MIT electronic circuit analysis program (EI). The MYCIN medical diagnosis project at Stanford [DI] has been constructing models of the techniques clinical pathologists presumably apply when attempting to make sense out of the raw data which describes some possible pathology or set of pathologies. In this sense they too are symbolically modeling cause-effect mechanisms. EL, Sussman's and Stallman's electronic circuit analysis program ([SS1] and [SS2]), is both an application program and a theory of circuit design. The main concerns of the project are: how are laws of electronics expressed in a way that is of use both to the analysis of a given circuit; e.g. the user specifies the starting states of all "active" components such as transistors, then asks the program to analyze the circuit's behavior, or the user requests the system itself to derive the various modes of operation of the active devices in the circuit and to the design of a circuit from descriptions of what it should accomplish.

## 4. Theoretical Framework

We have approached the representation and simulation of mechanisms from within the framework of our larger Commonsense Algorithms (CSA) Project. Since one main purpose of the project has been to understand more about how humans might represent and use all sorts of causal knowledge, we have explicitly chosen to develop the mechanisms theory in the context of the larger CSA theory, drawing upon existing CSA mechanisms which themselves seem to be broadly applicable in all aspects of cognitive modeling. Specifically, this
has resulted in (1) adopting as the mechanisms representation the same representation that is used in the other two phases of the project, namely, the plan synthesizer [R1] and the language comprehender [R1] and [R2], and (2) applying our notions of "spontaneous computation" [R3] to simulation. Although the running LISP Mechanisms Laboratory which has resulted is not ef $i$ iciency-wise competitive with applications simulators, this is of little concern, since efficiency considerations are not relevant to our current goals.

In summary, then, we want to be able to express and simulate mechanisms (and hence, also the physical principles upon which they are founded) in ways that might approximate how humans do these tasks. To accomplish this, we have chosen to embed the simulator in the same framework as the plan synthesis and language comprehension components of our project.

### 4.1. CSA Cause-Effect Representation

The declarative representation of mechanism is a cause-effect graph whose nodes are events and whose links are drawn from a set of the ten theoretical forms of inter-event causal interaction. Each event falls into one of four categories: action, tendency, state or statechange. With the exception of "tendency", these terms are intended to reflect their standard connotations. We use the term "tendency" to describe action-like events in which there is an absence of intention. Gravity, for example, is a tendency, since it is a force generator, but has no choice about, or reasons for, acting.

The links of our theory are intended to reflect what we believe to be necessary (and close to sufficient) underlying human concepts relating to causality in physical mechanisms. (These same links are also the basis of the plan synthesizer's representation.) We arrived at several of the ten by obvious intuitive reasoning, and at the remaining ones by considering many mechanisms and attempting to capture the recurring ideas. None of the links by itself is particularly novel to us, or individually provocative; yet, taken as a set, we believe these links are both theoretically significant, and habitable in practice.

### 4.2. Representation Example: Home Gas Forced-Air Furnace

We introduce the CSA mechanisms representation first via an example that illustrates both the individual links and the general size and complexity of mechanisms we are currently representing and simulating. The example is the Home Gas Forced-Air Furnace, and reflects our understanding of this relatively sophisticated device.

For the purpose of convenient presentation, we have segmented the description into three boxes: (1) thermostatic control, (2) heat generation, and (3) heat delivery. Interconnections $A$ through $G$ between boxes are indicated in circles. We have provided an event-wise cross-indexed English description of Box 1 to accompany its schematic representation, but have omitted the English descriptions of Boxes 2 and 3 for space reasons (see [RGl] for more details).

## BOX 1 (Control Subsystem)

We assume that the system heats a single room, and that this room contains a mercury-filled, glass-envelope style thermostat. A temperature change of the thermostat (1) is equivalent to a state of temperature fluctuation (2) (upper left hand corner). Such fluctuation continuously enables the tendency THERMAL EXPANSION (3) to produce a continuous change in the length of the (coiled) thermostat strip (4). Provided the glass envelope is attached to this coil (5), this length change is equivalent to a change in the angle (phi) of the glass envelope (6). The angle of the glass envelope can also be influenced by an external adjusting action (7,8). At any moment, the net change (9) of this angle is governed by these two causal sources.

As phi changes, there are two points of interest: one when phi is below zero (10), i.e. the envelope begins tilting to the left (11), and one when it is above at zero (16), i.e. begins tilting to the right (17). When the envelope is tilting left (11), and the left edge of the mercury is not already at the extreme left end of the envelope (12), the mercury is in a condition of being unsupported to the left (13). This allows the tendency GRAVITY (14) to manifest itself and to begin continuously changing the mercury's location toward the left (15). Returning to the other threshold point for phi (when phi is above zero), a symmetric system of unsupportedness pertains

( $16,17,18,19,14,21$ ), influencing the mercury to move toward the right. At any moment, the net change in the mercury's location (22) is governed by these two systems.

There are five points of interest (23,24,28,29,32) along the mercury's path of travel (lower part of Box l). When the left edge of the mercury reaches point $A$ (less than or equal to -2 ) (23), the mercury will cease being unsupported, and GRAVITY's influence will be severed (i.e. the mercury will stop moving left). While the mercury is between points $B$ and $C$ (i.e. between -3 and 1) (24) there will be physical contact of the mercury and the electrical contact pins P1 and P2 (25). This amounts to the thermostat having closed the furnace's control circuit (33). This condition feeds into Box 2 via tie point $A$, serving as one precondition for the main supply gas valve opening. Conversely, contact between the mercury and P1 and P2 ceases at points less than -3 or greater than 1 .

The Control Subsystem participates in a large feedback loop via tie point F from Box 3.

The descriptions of Boxes 2 and 3 are similar, and Box 1 introduces all the link concepts in the representation.



### 4.3. Mechanisms Cause-Effect Links

With this example in mind, we now give a very brief description of the ten links. See [RG1] and [R2] for more detailed discussions.

Continuous and One-Shot Causal Action $A$ or tendency $T$ causes state $S$ or statechange
 SC to exist, providing gating conditions Sl,..., Sn are in effect. For the continuous form, the action's continued presence is required to sustain the state or statechange (i.e. there is some other unspecified force which would annihilate $S$ or $S C$ if $A$ or $T$ were
 removed). For the one-shot form, $A$ or $T$ is required only momentarily. The gates sovern the effects the action or tendency will have on its environment, in that their continued and simultaneous presence during the action is required for the indicated causality "to flow" from the action or tendency to the state or statechange. These gates are distinct from A or $\mathrm{T}^{\prime}$ s enablements; enablements govern the execution of the action itself, independently from any effects it may produce.

## Continuous and One-shot Enablement



State $S$ enables action $A$ or tendency $T$. For the continuous form, $S^{\prime} s$ continued presence is required in order to begin and sustain $A$ or $T$. For the one-shot form, $S^{\prime}$ s presence is required only momentarily to allow the action to begin. (Semantically, one-shot enablement occurs when the performance of $A$ or $T$

modifies its environment in a way which liberates the action from the influence of the original enabling condition.)

## Continuous and One-Shot State Coupling

State S 1 or statechange SCl indirectly produces S 2 or
SC2. For the continuous form, $\mathrm{S}_{2}$ 's or $\mathrm{SC} 2^{\prime} \mathrm{s}$ continued existence is coupled to $\mathrm{Sl}^{\prime} \mathrm{s}$ or $\mathrm{SCl}^{\prime} \mathrm{s}$ continued existence. For the one-shot form, S 2 or SC 2 is independent of $S 1$ or $S C l$ after the initial coupling.


These links provide a means of expressing implicit
intervening causal relations which are either unknown or irrelevant to some description (i.e. states do not directly cause states). They accept gating states in the same manner as the causal links.

## State Equivalence

State Sl or statechange SCl is an eqvivalent formulation of state S 2 or statechange SC 2 , providing gating conditions $S G 1, \ldots, S G n$ are present. "Equivalent formulation" means that the two states are paraphrases-syntactically different expressions of the same underlying event. Equivalences, we feel, will always be present in any representation, and hence

ought to be dealt with explicitly. Most frequently, an equivalence plays the role of buffering two different points of view of the same event (e.g., the output of one mechanism is "photons exist", the input to another mechanism is "light present", and say, during mechanisms invention, we wish to join the two mechanisms). The existence or non-existence of either equivalenced event implies the existence or non-existence of the other.

## State Antagonism

State Sl or statechange SCl is antagonistic to state S2 or statechange SC2 (i.e. the two events are mutually exclusive), providing gating conditions SGl,...,SGn are in effect. This link is the companion of the state equivalence link.


## Rate Confluence

Statechanges $\mathrm{SCl}, \ldots, \mathrm{SCn}$ represent the culminations of multiple causal sources for a net statechange of some entity with respect to some feature. The net statechange, $S C$, specifies the composite rate as a symbolic expression which can be dynamically evaluated during a simulation. Syntactically, all contributory

statechanges, and the net statechange must reference the same entity, and the same varying feature of that entity (e.g. change in temperature of the heat exchanger).

## Threshold

Net statechange, $S C$, reaches a threshold, $S$, of interest to the description of the mechanism. $S$ is an instantaneous description of an entity with respect to the feature which is varying in the statechange, e.g. the temperature of the heat exchanger is 400 degrees. We distinguish positive and negative thresholds
 graphically and in the internal representation and simulation so as to provide for hysteresis.

This concludes the discussion of the CSA cause-effect links relating to mechanisms description. For interested readers, there are several other links relating to motivation and intentionality of human actors. These other links permit us to explore the areas of plan synthesis and language comprehension in the social as well as physical domain within the same representational framework.

We have represented a number of other physical, electrical and electronic mechanisms in these same terms, including (1) a computer flip-flop, (2) the "drinking duck" novelty toy, (3) an incandescent light bulb, (4) a reverse-trap flush toilet, (5) a mechanical oscillator, and (6) descriptions of composite events involving physical laws such as the Bernoulli effect, gravity and momentum. Additionally, we have used the representation to describe a computer algorithm for computing the average of a table of integers, and feel that the CSA links will provide a good language for programming concepts as well as physical laws. These and other examples appear in [RG1], [RI] and [R4].

## 5. Simulation Strategy

As we pointed out earlier, our purpose has not been merely to build a mechanisms description and simulation laboratory, but to build one upon concepts which are generally applicable to a wide spectrum of other types of cognitive modeling as well. We feel that we have succeeded in the description aspect by embedding the mechanisms description language in the same framework as our problem solving and language comprehension languages, namely these CSA
links just covered.

For the simulation strategy, we have drawn upon another aspect of the CSA system that we call the "spontaneous computation" component. This is a generalized implementation of pattern-directed inference, wherein computations occur spontanecusly, rather than on demand from another computation.

The simulation strategy (see flow diagram heiow) is this: convert the declarative (CSA) cause-effect representation of the mechanism to a population of autonomous computation units, each of which models one event in the declarative representation. If each unit in this population of spontaneous computations ( $S^{\prime} \mathrm{s}^{\prime}$ ) contairs a model of all other event schemata that can influence it, (both within, and external to, the particular mechanism in which the event participates) then the modeling $S C$ can be caused to run when all influences are just right by including all those influences in its (potentially complex) invocation pattern. Thus, our strategy is to model each declarative event by an $S C$ whose trigger pattern is sensitive to other prerequisite events in the mechanism.



#### Abstract

A simulation will amount to embedding the reactive $S C$ population which models the events in the mechanism in an environment in which static conditions are appropriate for triggering the mechanism. The environment, as well as all instantaneous states during the simulation, are modeled as a collection of database assertions. Triggering of the mechanism will lead to a conceptually parallel avalanche of activity wherein the running of one or more SC units can prompt the running of one or more other units, and so on. The simulation becomes quiescent when no remaining SC's perceive themselves to be relevant.

There are some important issues relating to why we have chosen this strategy for simtilation rather than more "straightforward" strategy that manipulates the declarative representation directly. We will return to the philosophy behind this approach to simulation later. We now describe the CSA system's spontaneous computation component which provides the substrate for the simulator.


### 5.1. Spontaneous Computation

Our implementation of spontaneous computation is a generalization of the pattern-directed invocation notions embodied in PLANNER [SWCI] and CONNIVER [MSI]. By "generalization", we mean specifically that we have provided for the specification and organization of more complex invocation patterns, and for more complex hierarchical organizations of $S C$ populations.

The invocation (trigger) pattern of an SC in the CSA system is constructed from nested n-tuples composed in virtually any degree of complexity using the logical relations AND, OR and ANY. Each component of the trigger can be one of the following types: (1) associative, (2) non-associative, or (3) computable. Associative trigger components come to be organized into a reactive data structure we call a "trigger tree"; these are the components which can react to passing stimuli (to be defined) and cause the entire pattern to be tested (polled). Non-associative trigger components represent aspects of the trigger's environment that must be true (when explicitly polled) in order for the SC to fire fully, but which themselves are incapable of initiating the firing. Computables are any EVALable LISP forms
other than (1) or (2), and must evaluate non-NLL for full triggering of the pattern to occur.

To illustrate, suppose we wish to create an $S C$ which will fire (i.e. spontaneously execute) whenever "the temperature of the heat exchanger is greater than 400 degrees in coincidence with either valve A or valve $B$ being open". We would express this condition by "planting" the following pattern in a reactive trigger tree:

$$
\begin{aligned}
& \text { (\$PLANT ' (AND (+ } 1 \text { (TEMP XCHGR -X)) } \\
& \text { (GREATERP -X 400) } \\
& \text { (OR } \left.\left\{\begin{array}{l}
-1(\text { OPEN VALVE-A) }) \\
-1 \quad(O P E N ~ V A L V E-B)
\end{array}\right)\right) \\
& \text { <some body> } \\
& \text { <some tree>) }
\end{aligned}
$$

The " + " form denotes an associative component, the " - " form denotes a non-associative component, and the GREATERP is a computable denoting itself. The integer l's indicate how much effort (in terms of database fetches by the deductive component) can be expended in the process of determining the presence or absence of all remaining parts of the pattern whenever the pattern is triggered via an associative component. Variables are prefixed by hyphen signs and are global to the entire trigger pattern, in that variables with the same name must be consistently bound across the entire pattern.

An SC body is an arbitrary EVALable LISP expression. Each SC is context sensitive, in that it can be masked and unmasked independently within its trigger tree. (The simulator relies on this feature to prevent looping in certain representation patterns.) Entire trigger trees are also context sensitive, as will be described.

The population of $\mathrm{SC}^{\prime}$ s represented by some trigger tree is caused to react to some stimulus (a fully-constant nested n-tuple) in either of the following ways:

```
                ($ACTIVATE <trigger tree> <stimulus>)
or (<trigger tree> <stimulus>)
```

where, in the latter form, the tree is used semantically as a function. Application of a trigger tree to a stimulus yields a queue of $S C$ 's which are determined to have been fully triggered, i.e., initially triggered through one
component, then proven fully applicable by calls on database/deductive system for remaining components.

### 5.2. Channels

Beyond this organization of complex trigger patterns into trigger trees, the CSA SC component provides for the higher-level hierarchical organization of trigger trees around constructions called "channels". An SC channel is intended to be the analog of a hardware channel, and a generalization of the PLANNER/CONNIVER pattern-directed invocation scheme.

A typical channel is depicted below, accompanied by the calls on the channel-constructing function, SCONNECT, which would set it up. A channel has one-dimensional extent; signals can be injected at arbitrary points, to propagate in a specified direction, and these signals constitute the stimuli to which trigger trees can react.


Two types of entities can be attached to the channel at "tap points", of which there may be virtually any number. Attached entities are either trigger trees ("watchers") or arbitrary LISP functions ("servers").

The channel construct itself amounts to a fracturing of what would otherwise be a private inter-function calling sequence. If, wherever function Fl would ordinarily call function F2, we now arrange to have Fl infect a request to $F 2$ over a channel (i.e. place the calling arguments as a signal on
the channel) other entities will then have a chance to inspect the signal as it passes, and (1) simply react benignly, (2) alter the signal as it passes, or (3) block the signal altogether. In this way, all function calls can be "fishbowled" by other servers and trees of spontaneous computations.

There may be any number of channels, and any given server or trigger tree can be attached to numerous channels, or to the same chansel at numerous tap points. The $S C$ 's in one trigger tree attached to one channel may, when they run, inject signals on other channels, and so forth, providing for cenmex hierarchical organizations of spontaneous computations into conce\%ual populations. The connection of an entity to a channel is context sensitive. and entities are disconnected via the function (\$DISCONNECT <entity> <channel>).

Signals are injected onto a channel via the function:
(\$INJECT <signal> <to-server> <on-channel> <in-relation> <to-tap-point> <moving-in-direction>)
i.e., inject a signal to some server on some channel at the injection point indicated by <in-relation> and <to-tap-point>, propagating in some direction. Thus, wherever F1 used to call F2 directly, when participating in the channel facility, it will now inject signals to $F 2$ on some channel on which F2 exists as a server, e.g.:
(\$INJECT: (TEMP XCHGR 400)' ${ }^{\prime}$ 'AT LETORE 'CHANNEL 1

That is, store a fact by injecting it to SSTORE of CHANNEL1, starting at the LEFTEND, propagating RIGHT. As the signal moves past watchers, or as the response signal back from $\$$ STORE moves past response-watchers, numerous $\mathrm{SC}^{\prime}$ s may be triggered and run.

The $S C^{\prime}$ s that the mechanisms simulator creates as the procedural representation of a mechanism all exist in two trees attached to two channels. The passing signals on these channels will be the symbolic descriptions of changing events in the simulation, entering and exiting the database via these channels. When it runs, a fully-triggered SC will place a new event description on one of these channels, and possibly mask or unmask itself
and/or other $\mathrm{SC}^{\prime} \mathrm{s}$ in the population (to be described).
We feel the trigger tree/channel paradigm is a powerful one. Currently, the mechanisms simulator employs only these two relatively simple channels in performing its tasks, (analogous to the store and erase channels of PLANNER and CONNIVER) but we expect it to utilize the full potentials of this system in future phases of the project, especially when we address the topic of mechanisms invention.

This concludes our brief sketch of the CSA SC component insofar as it relates to our purposes here. For a more complete discussion, including several other theoretical applications of the SC system in the areas of inference, problem solving and language comprehension, see [R3].

## 6. The Mechanisms Simulator

The simulation subsystem is a collection of LISP functions that take the internal declarative representation of mechanism, convert it to a population of spontaneous computations, then awaken a subset of the population via a triggering assertion. A subset of the awakened SC's will request that their bodies be run. The evaluation of the bodies will cause other SC's to awakened. This will constitute an execution of the mechanism. In the execution, events will be automatically asserted and deasserted in the database and states which are changing with time will be updated. The status of an event is established in accordance with its causal relationships to other events as they become asserted and deasserted.

The initiation of the simulator is caused by the assertion of some event: semantically, either an action performed by an external actor, a tendency relating to a natural force, or a state, as derived, e.g., as an output from the simulation of another mechanism. This latter initiating event will provide for mechanisms interaction where several mechanisms have a common event and the assertion of that event in one mechanism will initiate (or contribute to the initiation of) execution of the other mechanism.

Each link in the internal declarative representation of a mechanism has a set of $S^{\prime}$ 's automatically created for it according to a set of rules for each
theoretical link type. An $S C$ is activated when its preconditions have been satisfied. The running of an SC will generally cause a database modification or a masking or unmasking of another SC.

### 6.1. Declarative to Procedural Conversion

We now briefly give two examples to illustrate the process of declarative-procedural conversion. The SC's are diplayed with the following conventions. Each SC representing some potential event is imagined to be "eyeballing" a path to the database (channel). The path type (STORE or ERASE) is located between the double bars on the left side of the display. The logical connectives in the created SC's trigger pattecn are to the right of the triggering events and the arrows out of the triggering events flow into these connectors. On the far right side of the figure is the body of the SC. These are the activities which will be performed when the $S C$ is run. A reference name for the $S C$ is located on the arrow into the body of the $S C$.


To illustrate, in the example of the $S C^{\prime}$ s created for the continuous causal (C-CAUSE) link governing the movement of the mercury in the furnace's thermostat, when the tendency GRAVITY or the state "mercury unsupported on the right" is asserted and the other event is currently in the database (i.e., also present), then the state-change of the location of the mercury is stored
(i.e. asserted) in the database. If either the tendency GRAVITY or "mercury unsupported on the right" is deasserted (i.e erased) from the database, then the state-change of the location of the mercury is erased from the database. ()


The second example is the $S C^{\prime}$ s created for the continuous coupling (C-COUPLE) link which watches for the mercury to enter the region between Pl and P2. When the location of the mercury is between -3 and 1 there is contact between the mercury, P1, and P2 (i.e the contact event is asserted in the database). When the location of the mercury moves above or below the given range, the contact event is erased from the database.

This illustrates the conversion process on two of the simpler link types. More complex processing is required for other links. For example, the rate confluence ( $R$-CONFL) link requires symbolic rate computations to be performed. Details of the conversion process appear in [RGl].

### 6.2. Simulation Example

We have successfully simulated the entire furnace, illustrated earlier in its declarative form, although the simulation must be done in three pieces (the whole simulation will not fit into our limited memory). The furnace does indeed cut on, heat the room up, then cut off, providing a trace of all events in the sequence. To convey a picture of the simulator's operation, we include an excerpted and annotated sequence showing the simulation of the thermostatic control portion of the furnace.

The Mechanisms Lab requires that each mechanism's CSA description be coded into a machine readable form. This form is the external declarative CSA representation and is used as input to a mechanisms defining function called \$MECHANISM. This function generates the internal declarative representation of the mechanism from the external declarative representation. The call on \$MECHANISM with the appropriate external declarative form for the thermostatic control section of the furnace is shown below.

(\$MECHANLSM ${ }^{\circ}$ (
(NAME BOXI)
(NAME BOXI)
(EVAME BOX
(LINKS

(S-EQUIV $\left.\left(\begin{array}{ll}1 & 2\end{array}\right)\right)$
$\left(\begin{array}{l}\text { C-ENABLE }\end{array}\right.$
$\left(\begin{array}{ll}2 & 3\end{array}\right)$
$\left\{\begin{array}{l}\text { C-ENABLE } \\ \text { C-CAUSE } \\ \text { S-EQUIV } \\ \left.\left(\begin{array}{ll}2 & 4\end{array}\right)^{3}\right) \text { ) }\end{array}\right.$
$\left.\left\{\begin{array}{l}\text { S-EQUIV } \\ \text { C-CAUSE } \\ \left(\begin{array}{ll}4 & 6 \\ 7 & 8\end{array}\right)\end{array}\right)\binom{5}{5}\right\}$
(RATE-CONFL $\left(\binom{6}{8}\right.$ ) )

(S-EQUIV $\left.\left(\begin{array}{lll}1 & 2 & 1\end{array} 3\right)\left(\begin{array}{ll}1\end{array}\right)\right)$
(C-CAUSE $(1415)(13))$
(ANTAG (12 26))
$\left.\begin{array}{lll}\text { C-COUPLE } & (23 & 26\end{array}\right)$ )
RATE-CONFL ( 11521 )
THRESH $\left(22 \quad\left(\begin{array}{llll}15 & 21) & 22))\end{array}\right)\right.$

$\left.\begin{array}{lll}\text { C-COUPLE } \\ \text { C-COUPLE }\end{array}\left(\begin{array}{ll}28 & 12 \\ 29 & 18\end{array}\right)\right\}$
(C-COUPLE $(2425)$ )
$\left\{\begin{array}{l}\text { S-EQUIV }(2533) \\ \text { ANTAG } \\ \text { (27 }\end{array}\right.$
$\left\{\begin{array}{l}\text { ANTAG }(2718)) \\ \text { C-CAUSE }(1421)\end{array}\right.$
C-COUPLE $(1621)$ (19))
(C-COUPLE $\left(\begin{array}{ll}16 & 17))\end{array}\right.$
S-EQUIV (18 19) (17))
$\left\{\begin{array}{l}\text { THRESH } \\ \left.\text { THRESH }\left(\begin{array}{l}4 \\ \text { THR } \\ 1 \\ 1004)) \\ 1001)\end{array}\right)\right)\end{array}\right.$
(INITIAL-WORLD $14 \quad 5 \quad 200 \quad 201 \quad 202$ 203)
(TRIGGER 1) RATES
(RATEI RATE59)
(RATE:2 (\#QUOTIENT RATEI \#2\#0))
(RATE3 RATE2)
(RATE4 \#1)
(RATE 5 (\#PLUS RATE3 RATE4))
RATE6
RATE
RA )
(RATE7 \#4)
(RATE8 (\#PLUS RATE7 RATE6))
(RATE59 \#1) )
))

## Simulation Trace

The following illustrates the output of the simulator. There are 72 SC 's created for this section of the furnace. For space reasons, only a small portion of the trace has been left intact to convey the flavor of the simulation. For clarity, the edited trace has all database activities removed except those which affect the location of the mercury and those events directly linked to them. The left-hand side of the trace is an English description of the simulation.

## BEST AVALLABIE COPY



## 7. Simulation Philosophy

Why do simulation this way? Namely, why do it symbolically, and why go through the contortions of converting the declarative form to spontaneous computation units? Clearly, it would be possible to simulate a mechanism by applying relatively simple graph algorithms directly to the declarative representation.

To understand why we have adopted this SC strategy, reflect on the nature of any physical system - a mechanism in particular. A mechanism is built from physical principles and components that act autonomously, in the sense that they are governed by physical laws that "run in parallel". The mechanism just happens to work in desired ways because the inventor has managed to identify, harness, and coordinate a population of autonomous agents. In this setting, very ninor or very local alterations to one component or its environment can propagate to all parts of the mechanism in a falling-dominoes fashion. In particular, embedding the mechanism in alien or novel environments can significantly alter the micro and macro behavior of the mechanism in ways that will always relate to the individual components of the mechanism, viewed as autonomous agents, but never to the mechanism as a whole. Certainly, the net effect of embedding the mechanism in an environment will be to change the cumulative, overt behavior of the mechanism; but the overt behavior changes will be nothing more that the sum of many smaller, possibly unrelated influences.

By converting the mechanism's declarative form into SC form, we effectively crack the description open, exposing all the individual causereffect relationships directly to the environment. Each is free to behave as its trigger pattern dictates. Thus, when we embed the eviscerated form in a new environment, each cause-effect relationship must fend and produce results on its own.

In particular, this makes possible the more detailed study or debugging of the mechanism when placed in an environment where there are other mechanisms, as occurs, say, when a new mechanism is invented out of existing mechanisms and physical laws. For example, suppose that two mechanisms, designed separately, are thrust into the same environment and made to coexist, as happens when the two become part of a larger mechanism. Without cracking;
each decsription apart and casting the resulting spontaneous units into a large population in which the two mechanism's boundaries are lost, it would not be apparent how the two might interact. But with such a strategy, event-wise crosstalk is more readily discovered.** A case in point was the much publicized glitch in Polaroid's new SX-70 camera, where the sensitive electronics were interfered with by noise spikes produced by the concurrent operation of the picture ejecting motor. The interaction was a simple, but obscure one, and was solved by clever timing of events to remove the concurrency.
** Forcing two independent mechanisms to coexist in one environment is the analogy of forcing brother subgoals to coexist in one environment in plan synthesis; in both domains, there is the possibility of unanticipated interactions. See [RLI].

### 7.1. A Related Issue: Mechanisms Abstraction

There is quite of ten a need to suppress much of the representation or simulation detail present in the system's full model of a given mechanism. There would be a need for suppresssion of detail, for example, when we wished to provide a very high-level overview of a mechanism to serve, say, as a black box in a larger invention effort, or as an introduction of the mechanism to a totally naive CAI user. Since the suppression of mechanism detail is analogous to the suppression of detail in text, we use the term "abstraction".

One of our current areas of interest is in automating the process of mechanisms abstraction. Note that we already have provisions in the representation for expressing abstracted forms: the state coupling link. This link allows us the freedom of direct state-state causality. The abstraction technique can therefore be one of syntactically replacing certain patterns in the detailed representation with state coupling links.

It appears that there will be only a small number of syntactic replacement rules. Two of the most obvious abstracting rules are depicted below. Such rules, applied transitively, could perform both minor and major abstractions. A major abstraction would amount to a simple coupling between
the input and output states of a mechanism, e.g. "A falling temperature causes the the rmostat to close."

Such abstracting would be applied to the declarative representation before simulation, so that the converted computation units would directly reflect the abstraction. Abstraction could be applied uniformly over an entire mechanism, or locally to excise parts which are not of immediate interest. Non-uniform abstracting, in which all uninteresting sections have been defocussed, would deliver practical advantages by allowing the simulator to devote more time to the relevant aspects of the simulation. This could be important if, for example, we suspect the audio stages of a radio are at fault, but know the power supply is functional. Before simulation, the radio's power supply description could be defocussed before conversion to simulation representation (i.e. population of $\mathrm{SC}^{\prime}$ s).


Two Defocussing Rules

In addition to simple syntactic defocussing, we will also be exploring more semantic forms. Semantic mechanisms abstraction would be useful if, for example, we wished to know about the gas furnace with respect particular aspect of its operation, such as "heat production". We might scan through the description, assessing each event according to its semantic relevance to the concept of heat production. Then, using the highly ranked events as milestone events, the syntactic abstraction procedure could be applied to yield an abstraction with respect to the desired point of view, i.e. one which retained only the milestone states.

## 8. Conclusions

We have presented a theory of mechanism description and simulation which
is based on tenets and processes we hypothesize to be applicable to other aspects of human cognition. We feel the theory of cause-effect representation is sound and expressive for a wide variety of mechanisms, and that it bears significance as a theory of human cause-effect knowledge representation.

Our next specific goals are threefold: (1) to make the acquisition of new mechanism patterns interactive, having the model prompt the user, and verify that the user's use of the representation coincides with the model's notions of semantic well-formedness, (2) to study the processes of mechanisms abstraction, and (3) to apply the existing CSA plan synthesizer to the task of mechanisms invention, involving the simulator in a debugging loop.

## Acknowledgments

We wish to thank the other members of the Maryland CSA group (Phil London, Mache Creeger, John Boose, Georgy Fekete and Shmuel Peleg) for their participation in numerous discussions and for their suggestions. We wish also to thank the Office of Naval Research for its support of this research under Grant N00014-76c-0477.

## 9. References

[BBI] Brown, J., and Burton, R., Multiple Representations of Knowledge for Tutorial Reasoning, in Representation and Understanding, D. Bobrow and A. Collins (eds.), Ac ademic Press, 1975
[D1] Davis, R., Applications of Meta Level Knowledge to the Construction, Maintenance and Use of Large Knowledge Bases, Doctoral Dissertation, Stanford A.I. Lab. AIM 283, 1976
[JW1] Johnston, R., and White, M., Simulation of an Artificial Heart System, 1971 Summer Computer Simulation Conference, 1971
[L1] London, P., Abstraction Mapping as a Self-organizing Scheme for Problem Solving Systems, Doctorial Dissertation Proposal, University of Maryland, 1976
[LBRI] Luetscher, J., Boyer, D., and Resneck, J., Control of the Renal Circulation and Kidney Function, 1972 Summer Computer Simulation Conference, 1972
[MS1] McDermott, D. and Sussman, G., The CONNIVER Reference Manual, M. . T. AT Memo 259 a, 1974
[R1] Rieger, C., An Organization of Knowledge for Problem Solving and Language Comprehension, Artificial Intelligence, vol. 7, no. 2, 1976
[R2] Rieger, C., The Representation and Selection of Commonsense Knowledge for Natural Language Comprehension, Proc. Georgetown University Linguistics Roundtable, 1976
[R3] Rieger, C., Spontaneous Computation in Cognitive Models, University of Maryland TR 459, 1976 (to appear in Cognitive Science)
[R4] Rieger, C., The Commonsense Algorithm as a Basis..., Proceedings TINLAP Workshop, M. I.T., 1975
[RG1] Rieger, C., and Grinberg, M., The Causal Representation and Simulation of Physical Mechanisms, University of Maryland, TR 495, 1976
[RJSI] Raines, J., Jaffrin, M., and Shapiro, A., A Computer Simulation of the Human Arterial System, 1971 Summer Computer Simulation Conference, 1971
[RLI] Rieger, C., and London, Pubgoal Protection and Unravelling during Plan Synthesis University of Maryland, TR 512, 1977
[SLI] Sahinkaya, Y., and Lee, Y., A digital Computer Simulation Model for a SCR DC to DC Voltage Converter, 1972 Summer Computer Simulation Conference, 1972
[SS1] Sussman, G., and Stallman, R. Heuristic Techniques in Computer Aided Circuit Analysis, M.I.T. AI Memo 328, 1975
[SS2] Sussman, G., and Stallman, R. Forward Reasoning and i)ependency-Directed Backtracking,in a System for Computer-Aided Circuit Analysis, M.I.T. AI Memo 380 , 1976
[SWCl] Sussman, G., Wi nograd, T., and Charniak, E., MICRO-PLANNFR Reference Manual, M. I.t. AI Memo 203 a , 1971
[WBL1] Wright, E., Blokland, G., and Lawrence, B., Simulation of a Natural Circulation Boiler Using, a Hybrid Computer, 1972 Summer Computer Simulation Conference, 1992
Off of Naval Research Branch Office, Boston 495 Summer St.
Boston, Mass. 02210
New York Area Office
715 Broadway-5th Floor
New York, N.Y. 10003
Mr. E. H. Gleissner
Naval Shiy R+D Center
Computation and Math Department Code 18
Bethesda, Maryland 20084
Capt. Grace M. Hopper NAICOM/MIS Planning Branch OP-916D
Off, Chf, of Naval Up.
Washington, D.C. 20350
Mr. Kin B. Thompson
Technical Director
Information Systems Div. OP-91T
Off., Chf. of Naval Op.
Washington, D.C. 20375
Naval Research Lab.
Technical Info. Division
Code 2627
Washington, D.C. 20375
Dr. A.L. Slafkosky
Scientific Advisor
Commandant, USMC
Code RD-1
Washington, D.C. 20380
National Security Agcy.
Attn: Dr. Maar
Fort Meade, Maryland 20755
Off. of Naval Research
Code 1021P
Arlington, Va. 22217

Asst. Chief for Tech. ONR Dept. of Navy Code 200
Arlington, Va. 22217
Off. of Naval Research
Information Sys. Program
Code 437
Arlington, Va. 22217
Off. of Naval Research
Code 455
Arlington, Va. 22217
Off. of Naval Research
Code 458
Arlington, Va. 22217
Defense Documentn. Cent.
Cameron Station
Alexandria, Va. 22314
Off. of Naval Research
Branch Office, Chicago 536 South Clark St.
Chicago, Ill. 60605
Off. of Naval Research
Branch Off., Pasadena 1030 East Green St. Pasadena, Calif. 91106

Naval Electron. Lab. Ctr. Adv. Software Tech. Div. Code 5200
San Diego, Calif. 92152

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)
REPAORT DOCUMENTATION PAGE

Approved for public release; distribution unlimited.
17. OISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)
18. SUPPLEMENTARY NOTES
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Cause-effect representation Mechanism Invention
Declarative-procedural correspondence
Simulation of Physical systems
Spontaneous computations
20. ABSTRACT (Continue on reverse side if necesaary and identify by block number)

Abstract: A theory of cause-effect representation is used to describe man-made mechanisms and natural laws. The representation, consisting of 10 link types that interconnect events into large declarative patterns, is illustrated on a relatively sophisticated device, the home gas forced air furnace. Next, a procedure and framework for translating the declarative description of a mechanism into a population of associatively triggerable computation untts is described. The associative, or procedural, form can then be used to perform a discrete cause-effect simulation of the device. The delcarative to procedural translation, including a simulation trace, is shown for the furnace. Topics
of mechanism abstraction and mechanism invention are discussed, and the entire Mechanisma Laboratory is placed in the larger perspective of our research into human problem solving.

Keywords: cause-effect representation, declarative-procedural correspondence, simulation of physical systems, spontaneous computation, mechanisms invention

