

Tapir: the Evolution of an Agent Control Language

Gary W. King
University of Massachusetts
140 Governor's Lane
Amherst, MA 01003
gwking@cs.umass.edu

Marc S. Atkin
University of Massachusetts
140 Governor's Lane
Amherst, MA 01003
atkin@cs.umass.edu

David L. Westbrook
University of Massachusetts
140 Governor's Lane
Amherst, MA 01003
westy@cs.umass.edu

ABSTRACT

Tapir is a general purpose, semi-declarative agent control language that extends and enhances the Hierarchical Agent Control (HAC) architecture [1]. Tapir incorporates the lessons learned from developing HAC and makes it easier and faster to create reusable and understandable actions. Tapir has been used in a battalion level war-game simulation, a robot simulator, a simulation of cellular dynamics and a simulation of rodent behavior. The language is built around constructs that define agents, sensors, actions, and messages. It has mechanisms for handling multiple agents, a flexible resource model, and multiple means for structuring concurrent actions. This paper provides an overview of HAC and its shortcomings and then explains how Tapir extends and improves upon it.

1. INTRODUCTION

Building on previous work in simulation and agent control [1, 7, 19] we have created Tapir: an expressive agent control language with a simple syntax. Tapir has been used in multiple systems by experienced and novice agent programmers and has fostered dramatic increases in productivity. Tapir constructs are self-contained and modular, making them human-understandable and machine-parsable. Tapir is a semi-declarative, general purpose agent control language implemented in Common Lisp. Tapir has constructs for defining agents, sensors, actions, and messages:

Agents Agents are viewed primarily as resources for actions. They connect actions to the physics of the real or simulated domain. As resources, agents come in many different types. For example, they can be serializably reusable, sharable, composite, consumable and so on. Individual resources may exist in a hierarchy. For example, a robot is an agent consisting of motor resources, camera resources and a gripper.

Sensors Sensors tie actions to the world. They can be “primitive”, connecting to the world directly, or they can be abstract, amalgamating and processing data from multiple sources. Sensors can be shared by multiple actions.

Actions Actions use resources and sensors to carry out their tasks. Each action exists in the context of a control hierarchy and may extend the hierarchy by creating children of its own. This hierarchy is grounded in “primitive” actions that use resources to achieve their ends. Tapir has many constructs for structuring control, including sequential, parallel and repeated execution.

Messages Actions, sensors, and resources communicate via messages. For example:

- sensors send actions messages when events in the world occur.
- child actions send their parent messages when they complete or when they need to communicate some change in status.
- parents send their children messages when they wish to stop or redirect their activities.
- resources send their actions messages when they change or are destroyed.

Each message is an instance of a particular class and can carry additional information. For example, a FAILURE message carries the reason for its failure and a CHANGE-SPEED message carries the value of the desired speed.

Tapir is built upon and extends the Hierarchical Agent Control (HAC) architecture [4, 1]. The rest of this paper provides a brief overview of HAC (section 2), describes Tapir’s main constructs (section 3), and compares Tapir with existing languages (section 4). We conclude by discussing current applications and our future work.

2. THE HAC ARCHITECTURE

HAC is a framework for controlling agent behavior. HAC takes care of the mechanics of executing the code that controls an agent, passing messages between actions, coordinating multiple agents, arbitrating resource conflicts between agents, and updating sensor values. HAC structures agent control into three hierarchies: control, sensor and context. These determine what to do, what is happening, and how to interpret it. The last is especially important in planning since the correct response to an event depends on the goals and interests of an action’s ancestors. HAC is action-centered, not agent-centered. Rather than specifying what an agent is to do, one specifies what is to be done. Actions then use whatever resources are available to accomplish their ends. We found that this inversion of the usual framework made it easier to incorporate both planning and multi-agent control into the architecture.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

| | | | | | |
|--|------------------------------------|-------------------------------------|----------------------------|---|---------------------------------|
| 1. REPORT DATE 2002 | | 2. REPORT TYPE | | 3. DATES COVERED 00-00-2002 to 00-00-2002 | |
| 4. TITLE AND SUBTITLE Tapir: the Evolution of an Agent Control Language | | | | 5a. CONTRACT NUMBER | |
| | | | | 5b. GRANT NUMBER | |
| | | | | 5c. PROGRAM ELEMENT NUMBER | |
| 6. AUTHOR(S) | | | | 5d. PROJECT NUMBER | |
| | | | | 5e. TASK NUMBER | |
| | | | | 5f. WORK UNIT NUMBER | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Massachusetts, Department of Computer Science, Amherst, MA, 01002 | | | | 8. PERFORMING ORGANIZATION REPORT NUMBER | |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) | | | | 10. SPONSOR/MONITOR'S ACRONYM(S) | |
| | | | | 11. SPONSOR/MONITOR'S REPORT NUMBER(S) | |
| 12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited | | | | | |
| 13. SUPPLEMENTARY NOTES | | | | | |
| 14. ABSTRACT | | | | | |
| 15. SUBJECT TERMS | | | | | |
| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT | 18. NUMBER OF PAGES 8 | 19a. NAME OF RESPONSIBLE PERSON |
| a. REPORT unclassified | b. ABSTRACT unclassified | c. THIS PAGE unclassified | | | |

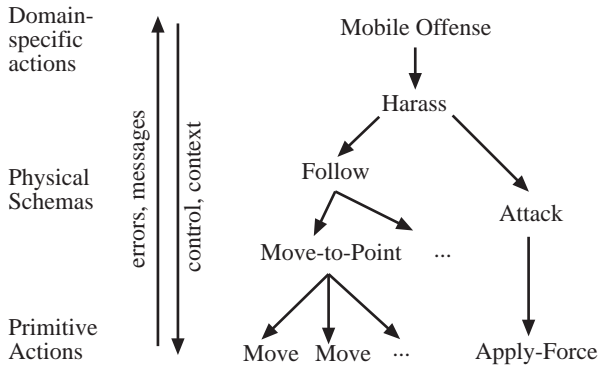


Figure 1: Actions form a hierarchy; control information is passed down, messages are passed up. The lowest level are agent effectors.

2.1 Agent control in HAC

HAC executes actions by scheduling them on a queue. The queue is sorted by the time at which the action will execute. Actions are taken off the queue and executed until no more actions are scheduled to run at the current time step. Actions can reschedule themselves, but in most cases, they will be rescheduled when woken up by messages from their children, sensors or resources. Action's execute their code during their **do** phase. More deliberative actions will tend to execute their **do** phase only intermittently whereas more reactive ones may execute theirs almost continuously.

HAC is a *supervient* architecture [18]. It abides by the principle that higher levels should provide goals and context for the lower levels, and lower levels provide sensory reports and messages to the higher levels (“goals down, knowledge up”). A higher level action cannot overrule the sensory information provided by a lower level, nor can a lower level interfere with the control of a higher level. Supervenience structures the abstraction process; it allows us to build modular, reusable actions. HAC simplifies this process further by enforcing that every action's implementation (its **do** phase) take the following form:

1. Respond to messages coming in from children.
2. Update state.
3. Schedule new child actions or use resources as effectors to alter the world.
4. Send messages up to parent.

Figure 1 shows a small part of an action hierarchy. The FOLLOW action, for example, relies on a MOVE-TO-POINT action to reach a specified location. MOVE-TO-POINT will send status reports to FOLLOW if necessary; at the very least a completion message (failure or success). The only responsibility of the FOLLOW action is to issue a new target location if the agent being followed moves. HAC is an architecture; other than enforcing a general form, it does not place any constraints on how actions are implemented. Every action can choose which messages it will respond to. Although actions lower in the hierarchy will tend to be more reactive, whereas those higher up tend to be more deliberative, the transition between them is smooth and completely up to the designer. Unlike other architectures [14, 7, 12], we do not prescribe a preset number of behavioral levels. Parents can run in parallel with their children or only when the child completes.




Figure 2: Raw sensor data is transformed into more complex concepts via the abstract sensor hierarchy.

2.2 The Sensor Hierarchy

The sensor hierarchy provides a principled means of structuring the complexity of reading and transforming sensor information. It functions analogously to the HAC action hierarchy and reduces the complexity of sensor fusion. The sensor hierarchy is grounded by the low level primitives available from the physics of the world (real or simulated). These primitive sensors include the location of terrain features, the current speed and location of agents in the world, the status of a robot's bump sensors and so forth.

Each level in the hierarchy integrates and extends the level below it by compiling the available information and providing additional structure. We call these higher levels *abstract sensors*, since they do not sense anything *directly* from the world. For example, enemy location information can be combined into a sensor that specifies overall enemy presence; terrain information can be combined into a sensor that specifies passes and movement corridors. Furthermore, these two sensors can be combined to show enemy vulnerability: areas where enemy units are concentrated and cannot move quickly (see Figure 2).

The sensor hierarchy shares the control hierarchy's syntax and structure. Each sensor is analogous to a HAC action. It sends and receives messages and performs sensor computation during its **do** phase. One advantage of this is that the same principles learned in building actions carry over directly when building sensors. This linkage also makes it easy for actions to use sensors as part of their control mechanism. An action can react to sensor messages the same way it reacts to those from child actions. Each sensor is associated with the set of actions that request it and completes when this set becomes empty.

Like actions, sensors abide by the principle of supervenience. Higher-level sensors integrate and interpret lower-level ones but they do not change the lower-level information. Lower-level sensors provide information to the higher-level ones but they do not tell them what to say. One advantage of this is that each level of the hierarchy can be viewed independently without worrying about the levels coming into it or the levels that are using it.

In summary, HAC provides a simple and flexible architecture for agent control and sensor fusion. Tapir extends the HAC framework by codifying standard idioms, supporting complex action constructs, simplifying resource use, adding declarative meta-information and making actions easier to build, modify and debug.

```

(defaction move-to-random-point ()
  :documentation "Move an agent to a randomly
    selected location."
  :resources (agent)
  :do
    (move-to-point
     :target-geom (find-random-location-for-agent
                   the-simulation agent))
  :on-message (message
              (stop-children)
              :restart))

;; -----
(defaction swarm ()
  :documentation "Move any number of agents to
    random locations repeatedly."
  :resources ((agents :count :all))
  :do
    (:foreach agent in agents :in-parallel
     (move-to-random-point :agent agent))
  :ignore-messages)

```

Figure 3: Implementation of multi-agent “swarm” behavior in Tapir.

3. TAPIR ARCHITECTURE

Tapir gains some of its power and simplicity by limiting the flexibility of the underlying HAC framework and clearly specifying the model that it implements. In particular, Tapir specifies a simple process model, a resource description language, and a limited but powerful set of control constructs. Tapir also adds support for debugging, data collection, interactive design and an extended syntax that helps clarify the action writer’s intent. Each of these will be discussed in turn after we first provide examples of some simple Tapir actions.

Actions are defined with **defaction**, a Lisp macro that defines a **CLOS** class [20, 15] and the methods that implement it. Actions can be used as building blocks for the creation of further actions. For example, figure 3 shows how **MOVE-TO-RANDOM-POINT** can be used to build **SWARM**. **SWARM**’s clauses specify that it should bind all of the resources available to it in its **agents** slot and create a **MOVE-TO-RANDOM-POINT** action for each of them in its **do** phase. The **:FOREACH** control construct is expanded at run-time and allows an unlimited number of children to be executed in parallel or sequentially. **MOVE-TO-RANDOM-POINT**’s clauses specify that it should bind only a single resource (the default) to its **agent** slot and that its **do** phase consists of starting a single **MOVE-TO-POINT** action with a random destination. Finally, the **:ON-MESSAGE** clause indicates that **MOVE-TO-RANDOM-POINT** should restart whenever it receives any kind of message.¹ This might be a message from the **MOVE-TO-POINT** child action or one from the agent it is using as a resource.²

3.1 Process Flow in Tapir

Figure 4 displays a schematic of an action’s life cycle. In brief: an action is created, enters its **do** phase (perhaps many times) and then completes. The details of this simplified description are filled in by special initialization and finalization code, sensors, process monitors, child actions, resource utilization and more. Figure 4 provides

¹Messages are typed and **MESSAGE** is the root superclass for all other message types.

²It is instructive to compare the Tapir version of **SWARM** with its HAC counterpart in [2]. The Tapir version is significantly simpler and shorter.

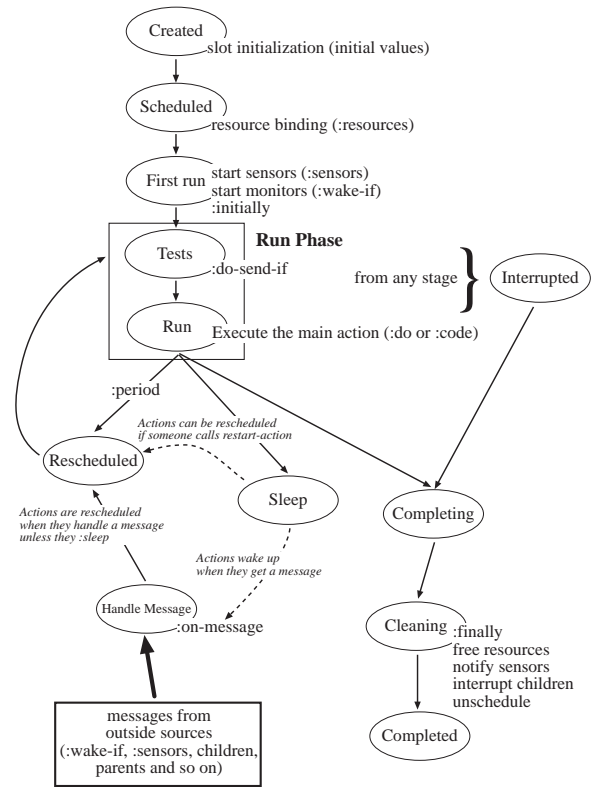


Figure 4: Process Flow in Tapir.

both a schematic of the process flow and the Tapir language constructs that can be used to influence and describe what an action should do. We summarize the phases below.

Initialization Actions maintain internal state in their slots; they also may require resources to do their job. When an action is created, its slots are initialized and any required resources are found and bound. Created actions do not do anything until they are actually scheduled to be run and the time for their activation arrives.

The Do Phase When activated, the action enters its **do** phase. Actions may execute special code at the beginning of the very first **do** phase; any sensors required by the action are also found or created at this time. The **do** phase itself consists of pre-checks to see if any messages should be sent to the action’s parent (this may possibly complete the action without it ever having executed); the actual execution; and then post-checks to see if any messages should be sent.

Sleeping Once an action executes, it will go to sleep unless it has been rescheduled. A sleeping action remains asleep until it is explicitly rescheduled (perhaps by a parent) or until it receives a message.

Message Handling Actions specify how to respond to each message type. These responses include: running code, re-entering the **do** phase, sending the message on up to the action’s parent, and ignoring the message completely.

Completion and Cleaning When an action is sleeping and it has no sensors or children that can wake it up, it automatically completes. Actions can also be interrupted at any time. In both cases, the action enters a cleaning phase where resources are unbound, finalization code is run, children are interrupted

and so on.

As mentioned above, messages are used to control the flow of execution between parents and their child actions. Children can complete themselves explicitly by sending a message of type completion (e.g., success or failure) to their parents. If a child completes for some other reason (e.g., a required resource is destroyed) then it will automatically send a completion message during its clean-up phase. Message passing is the only means of control transfer.




Figure 5: Tapir Resource Classes.

3.2 Resource control

Resources are the connections from an action to its domain. They can be agents or parts of agents. At the lowest level, “primitive” resources have values that alter the physics of the domain (real or simulated). For example, setting the wheel speed of a robot causes it to move, and setting the target-region of an artillery unit causes it to fire into that area. The Tapir resource model consists of two overlapping ontologies: one of individuals and groups, and one of parts and wholes (see figure 5). Depending on the domain, one or both of these ontologies may be more important. For example, in the Capture the Flag wargaming simulation [2] the individual agents are brigade and battalion sized military units (*wholes*). Some actions find it convenient to group these units into larger forces (*groups*). Furthermore, no unit can be used by more than one action at a time (the resources are serializably reusable). On the other hand, the agents in the robot domain are robots (*wholes* and *composites*) each with a movement system, vision system, gripper and other low-level effectors (*components*). Actions in this domain may require an entire robot or even a group of robots. It is also possible for an action to need only the wheels or only the gripper. Furthermore, it is possible for multiple actions to share parts of a robot at the same time.

Each domain therefore requires a different model for its agents, each with their own properties. Tapir uses the **defresource** macro to specify how the resources of a domain behave. This includes indicating whether a resource is primitive, its composite structure, whether it can be shared, whether it can be grouped, whether it is consumed and so on. Given a resource model, each action can then specify which resources it requires using its `:RESOURCES` clause. When a parent action starts a child, it can specify exactly how the child should bind its resources or it can simply specify which resources are available and leave the binding decisions to the child. Each resource specification in the `:RESOURCES` clause may include information about the type of the desired resource, the desired count and a predicate which each resource must satisfy. Tapir uses these specifications to validate parent specified bindings and as inputs to a greedy resource binding algorithm for child deter-

mined ones. Finding optimal resource bindings is a full fledged scheduling problem which Tapir does not attempt to solve.³

3.3 Child Control Constructs

An action’s `:DO` clause specifies the potential sub-tasks of an action. It describes both the sub-tasks and how they will be executed. The `:DO` specification is grounded in the sub-tasks which can be either Lisp code (using the `:CODE` construct) or an action specification consisting of the child action’s name and any of its parameters. The sub-tasks can be joined together using the following constructs:

:case This is analogous to the case or switch statements found in many languages. A form is evaluated and the result is used to determine which action specification to execute.

:foreach Dynamically creates a list of children and runs them. The children may be run using any of the child combination types (e.g., in parallel, in sequence and so on).

:in-parallel Runs a list of children in parallel. Each child in the list is started at the same time and runs independently. The order in which the children are initially started is unspecified.

:in-sequence Runs a list of children in sequence. Only one child at a time is active.

:repeat Runs a child action repeatedly for a count or until some predicate becomes true.

:one-of This is Tapir’s version of a the Lisp **cond** form. It takes a list of predicates and action specifications. **ONE-OF** runs the first action specification whose predicate evaluates to true. An optional `:OTHERWISE` clause can be included to run a specification when none of the other clauses is applicable.

:unordered Runs a list of children in some arbitrary sequence. Like **:IN-SEQUENCE**, only one child is active at a time but the order of activation is unknown.

:when Each when clause consists of a predicate and an action-specification. If the predicate evaluates to true, the action-specification is run.

Simple actions are easy to build because of Tapir’s straightforward syntax. The action combination mechanisms mean that it is also easy to build complex actions out of these building blocks. Control hierarchies that would have been difficult to build in raw HAC can be written easily in Tapir.

3.4 Miscellaneous

Tapir provides many facilities that speed action creation, debugging and use.

3.4.1 Superclass inference

Action behavior is specified in part by the HAC infrastructure classes from which it inherits. For example, there are classes for child actions, parent actions, agent using actions, agent managing actions, periodic actions and so on. Part of defining an action in HAC is determining its proper superclasses. This is often a non-trivial problem due to the number of superclasses and the interactions between

³Although the mechanisms which Tapir uses are available for use in a planner.

them. Tapir removes this problem entirely by automatically inferring the proper HAC superclasses based on the action specification. Tapir still supports action inheritance, however, which makes it possible to build generic action *pieces*. These pieces tend not to be complete actions but instead add functionality such as shared parameters or resource specifications.

3.4.2 Parameter checking and type coercion

Tapir's slot syntax (in the `:PARAMETER` and `:LOCAL` clauses) lets action writers explicitly describe slot properties that could be included only as comments in HAC. For example, slots can be marked as `:READ-ONLY` or `:REQUIRED`. Furthermore, writers can specify a slot's `:TYPE` and define type coercion functions. These last can be generally applicable or applicable to only a specific action. Coercions exist to convert lists of agents into a group agent, to convert 2-dimensional geometries into their locations and so on. By making the slot types explicit and providing a framework for principled coercion, Tapir makes the semantics of each action clearer and makes it easier for action users because they can let Tapir take care of the details.

3.4.3 Documentation and Consistency tools

Tapir's `:DOCUMENTATION` clause provides a place for writers to describe the intent of an action and its parameters. Tapir combines all of this documentation into a single form and makes it immediately accessible (via the Lisp documentation facility). Tapir also maintains a tree of sensors and actions and the sensors and actions that they use (in the `:SENSORS` and `:DO` clauses) along with the parameters passed to them. This tree allows Tapir to do compile-time consistency checks whenever a parent or child is defined or changed. Both of these tools make it easier to find the parameters of actions and sensors and to ensure that child actions and sensors are called correctly.

3.4.4 Interfacing with a simulation

Tapir is a general purpose agent control language and can be used in a myriad of ways. Because each simulator or real world environment has a different set of primitives and functions, Tapir includes a simple *translation* facility so that it can be customized to fit. This feature makes it easier to adapt Tapir to new environments and, perhaps more importantly, makes it easier for action writers. For example, we can translate simulator functions into new constructs that are more uniformly named, contain additional error checking and hide common parameters.

The facility itself capitalizes on Lisp's macro abilities and can parse elements of an action definition in arbitrary ways. For example, the translation for `stop-child` converts (`stop-child attack-action`) into

```
(when attack-action
  (interrupt-action the-simulation
                   attack-action child))
```

and (`:announce 'units-in-area`) converts into

```
(send-response-message-to-actions
 the-simulation the-action 'units-in-area)
```

The `THE-SIMULATION` and `THE-ACTION` arguments are translations themselves and expand into code that references the current simulation environment and the current action respectively.

3.4.5 Debugging and Instrumentation

Tapir's `:DEBUG` and `:INSTRUMENT` clauses make it easy to add debugging code to an action and to collect information while an action is running. Each `:DEBUG` clause creates switchable code or print statements. By default, the name of the debug switch is the same as that of the action although the `:DEBUG-NAME` clause makes it possible to change this. Debugging switches can be turned on and off with the `debug` and `undebug` commands. When the switch for an action is on, the debug code and print statements will be executed. This low-level facility makes it very easy to see exactly what is going right—or wrong—during an action's execution.

Action instrumentation relies on the EKSL CLIP package [21]. CLIP allows data to be collected from running systems by hooking metaphorical alligator clips onto variables just as one can clip wires onto an electric circuit. The `:INSTRUMENT` clause adds CLIP instrumentation to individual action variables (`:PARAMETERS` and `:LOCALS`) or to the action as a whole. The author can specify how often the CLIP's should be collected (for example periodically or only when a certain trigger condition obtain). This facility makes it easy to build hooks into a system for later experimentation and analysis.

3.5 An extended example

We conclude our discussion of Tapir's architecture with a substantive example (figure 7). The `MOVE-IN-FORMATION` action is used in the Capture the Flag wargaming simulator to move a group of agents in formation. Like `SWARM`, this action binds all of the resources available to it. It uses the `:INITIALLY` clause to determine which agent would make the best leader (using `select-leader`) and makes the rest of the agents followers. The `:DO` clause creates two parallel hierarchies—one for the leader and one for the followers. The former runs a leader move whereas the later determines formation criteria and then creates a separate move action for each follower and runs all of these in parallel. The action completes either when the leader move completes or if all of the agents are destroyed. The `:ON-MESSAGE` clause handles `resources-gone` messages from the agents of the action. A `resources-gone` message is generated when a resource is destroyed. The `MOVE-IN-FORMATION` action handles this message by selecting a new leader, stopping all of its children and restarting. A portion of the action hierarchy created by this action is shown in figure 6. For comparison, the original HAC code for `MOVE-IN-FORMATION` required more than three times the number of lines! Furthermore, the HAC code was spread out over four classes and seven methods.

4. RELATED WORK

Although motivated largely by its HAC underpinnings, Tapir shares much with other agent control architectures. For example, ESL [11], DAML-S [6] and PDDL [17, 16] roughly match Tapir's expressive power (see [5] for a comparison). But whereas ESL adds an action language to Common Lisp, Tapir wraps an action language around it. This difference is motivated in part by our desire to craft a simpler language usable by non-programmers but it also tends to make Tapir actions more modular and comprehensible.

PDDL and DAML-S are more concerned with describing services than actually implementing them. They tend towards the declarative whereas Tapir is more procedural in nature. We see these two poles as complementary and hope to weave some of DAML-S's ontologies into Tapir's executable model.

The APEX architecture also attempts to manage multiple tasks in

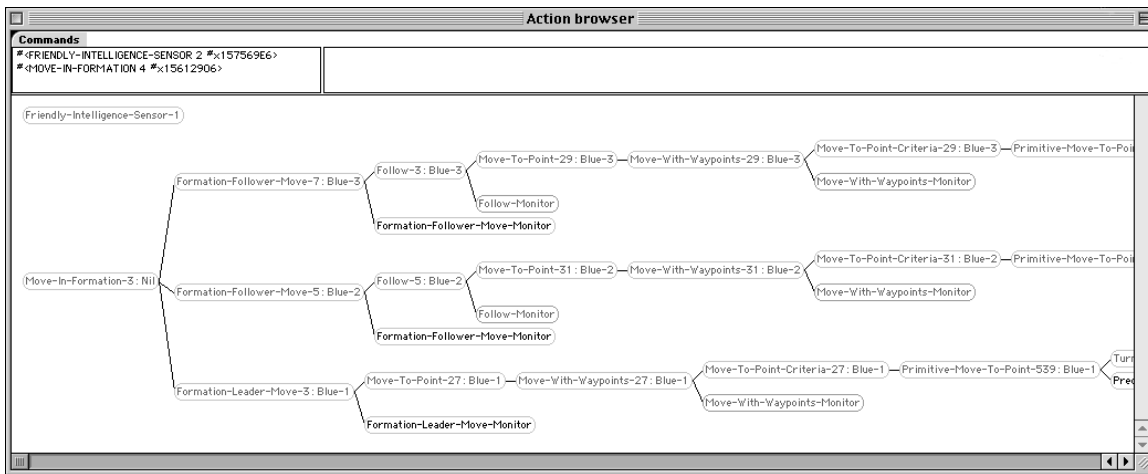


Figure 6: HAC Action Browser showing a portion of MOVE-IN-FORMATION’s control hierarchy

complex, uncertain environments, placing particular emphasis on the problem of resolving resources conflicts [10]. The current version of Tapir is less concerned with planning and scheduling and more concerned with letting authors tell agents what to do. Future versions of Tapir will add a planning language superset.

Like PRS [13], Tapir allows for the specification of blocking and non-blocking children (child actions that run in sequence with their parents or in parallel), and like later versions of RAP [8], success and failure are treated like any other message, and do not implicitly determine the flow of control between actions.

Tapir and HAC use the same representation for actions at all levels of the hierarchy, and also for sensors. Contrast this with the majority of current agent control architectures, e.g. CYPRESS [22] and RAP [9], which distinguish between procedural low-level “skills” or “behaviors” and higher level symbolic reasoning. Different systems are often used to implement each level (CYPRESS combines SIPE-2 and PRS, for example). Tapir does not conceptually differentiate between discrete actions and continuous processes, nor does it limit the the language used to describe them.

5. CONCLUSION AND FUTURE WORK

Tapir has achieved many of its initial design goals: it is a simple, flexible and intelligible agent control language that is easy to use and understand. Even so, Tapir remains a work in progress. We are investigating three extensions: planning, real robot control and better programming environments.

Part of the Tapir vision is to build a language clear enough for non-AI Subject Matter Experts (SMEs) to learn and use. Although Tapir is a relatively simple language it is not simple enough. There are at least three reasons for this:

1. Tapir is a programming language and, as such, is particular about its syntax,
2. Tapir is *only* an agent control language and does not have constructs for talking about the domain in which the agents function,
3. Agent control is inherently difficult.

We cannot do anything about the last item but we do hope to find partial solutions to the first two by extending the language so that it can talk about domains and by extending the Tapir programming environment so that it does more for the action writer. We call this project Visual Tapir.

5.1 Visual Tapir

Programming is replete with syntax, keywords, options, flags and clauses that are difficult for non-programmers to remember. Like other visual tools, Visual Tapir will add scaffolding to the programming environment so that authors can focus on their goals and not on minutia. Visual Tapir will be divided into a core environment and domain specific extensions. The former will focus on agent control proper whereas the later will add domain knowledge and functionality (for example, selecting locations on a map or specifying the features of a robot environment).

5.2 Tapir for planning

Tapir is an agent control language but it is not a planning language. We intend to extend Tapir with constructs for pre- and post-conditions, invariants, temporal reasoning and goal specification by merging it with HAC’s GRASP Planner [3]. Doing so will make Tapir significantly more powerful and more useful. The difficulty will be to keep it simple enough for SMEs to use productively.

5.3 Tapir for real robots

We have used Tapir to control simulated robots and our goal is to extend it (and the HAC substrate) so that actions can control real Pioneer II robots. The current HAC engine uses a centralized queue and imposes no constraints on the CPU time used by an action. Future engines will be operating in a real-time, decentralized environment, and will need to deal with widely varying time scales, from microseconds to days.

5.4 Conclusion

This paper has introduced Tapir, an agent control language built on top of the supervenient HAC architecture. Tapir highlights include:

- A simple and consistent syntax for the specification of actions, sensors, process monitors, message handlers, slots, and resources.

```

(defaction move-in-formation ()
  (:resources ((agent :count :all
                    :group-type group-blob)))
  (:parameters ((formation-shape :required)
                (distance 1.0)
                (target :required)))
  (:locals ((leader nil)
            (followers nil)
            (assignments nil)))
  (:initially
   (:code
    (setf leader (select-leader the-simulation
                              the-action)
          followers (remove leader (resources agent))))
  (:on-message
   (resources-gone
    (:debug "MIF: ~A died" (name leader))
    (setf leader (select-leader the-simulation
                              the-action)
          followers (remove leader (resources agent)))
    (stop-children)
    :restart))
  (:do
   (:debug "MIF: Moving to ~A with ~A leading"
          target leader)
   (:in-parallel
    (:in-sequence
     (formation-leader-move :agent leader
                           :target target
                           :followers followers)
     (:generate completion)))
   (:in-sequence
    (:code
     (setf assignments
       (assign-followers-to-criteria
        the-simulation leader followers
        (create-criteria formation-shape
         leader followers))))
    (:foreach (follower . assignment)
     in assignments :in-parallel
     (formation-follower-move
      :agent follower
      :target-agent leader
      :distance (distance assignment)
      :angle (angle assignment)
      :final-target target
      :use-follow-if-further-than
      (* 2.0 distance))))))
  :send-messages-to-parent)

```

Figure 7: Multi-agent formation movement in Tapir.

- A powerful action combination facility that makes it easy to build complex actions out of simpler building blocks.
- A flexible resource ontology, description language, and specification language.
- A set of helpful tools include debugging aids, syntax customization, and documentation facilities.

Tapir has been used to control units in a simulated war game, simulated robots and cellular dynamics. No controlled studies have been run to validate its utility but informal evidence indicates that Tapir users enjoy the language and are significantly more productive.⁴

Acknowledgments

This research is supported by DARPA/USAF under contract numbers F30602-01-1-0589, N66001-00-C-801/34-000TBD, DASG60-99-C-0074, F30602-99-C-0061, and F30602-00-1-0529. The U.S.

⁴Using the admittedly imprecise Lines of Code metric, Tapir actions require from two to four times *less* effort.

Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements either expressed or implied, of the Defense Advanced Research Projects Agency/Air Force Materiel Command or the U.S. Government.

6. ADDITIONAL AUTHORS

Additional authors: Paul R. Cohen (University of Massachusetts, email: cohen@cs.umass.edu)

7. REFERENCES

- [1] M. Atkin, D. Westbrook, and P. Cohen. HAC: A unified view of reactive and deliberative activity. In *Working Notes of Fourteenth European Conference on AI Workshop on Balancing Reactivity and Social Deliberation in Multi-Agent Systems*. ECAI, 2000.
- [2] M. S. Atkin, G. King, D. Westbrook, B. Heeringa, A. Hannon, and P. R. Cohen. SPT: Hierarchical agent control: A framework for defining agent behavior. In *Proceedings of the Fifth International Conference on Autonomous Agents*. Autonomous Agents, 2000.
- [3] M. S. Atkin, D. L. Westbrook, and P. R. Cohen. Domain-general simulation and planning with physical schemas. In *Proceedings of the 2000 Winter Simulation Conference*, pages 1730–1738, 2000.
- [4] M. S. Atkin, D. L. Westbrook, P. R. Cohen, and G. D. Jorstad. AFS and HAC: Domain-general agent simulation and control. In *Working Notes of the Workshop on Software Tools for Developing Agents, AAAI-98*, pages 89–95, 1998.
- [5] J. Blythe. Mappings between SADL and other action languages (available at <http://www.isi.edu/expect/rkf/sadl-mapping.html>). Technical report, University of Southern California, November 2001.
- [6] M. Burstein, J. Hobbs, O. Lassila, D. Martin, S. McIlraith, S. Narayanan, M. Paolucci, T. Payne, K. Sycara, and H. Zeng. DAML-S 0.5 draft release (available at <http://www.daml.org/services/daml-s/2001/05/>). Technical report, DARPA, May 2001.
- [7] P. R. Cohen, M. L. Greenberg, D. M. Hart, and A. E. Howe. Trial by fire: Understanding the design requirements for agents in complex environments. *AI Magazine*, 10(3):32–48, Fall 1989. also Technical Report, COINS Dept, University of Massachusetts.
- [8] R. J. Firby. Task networks for controlling continuous processes. In *Proceedings of the Second International Conference on Artificial Intelligence Planning Systems*, pages 49–54, 1994.
- [9] R. J. Firby. Modularity issues in reactive planning. In *Proceedings of the Third International Conference on Artificial Intelligence Planning Systems*, pages 78–85, 1996.
- [10] M. Freed. Managing multiple tasks in complex, dynamic environments. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pages 921–927, Madison, WI, 1998.

- [11] E. Gat. ESL: a language for supporting robust plan execution in embedded autonomous agents. In *Proceedings of the IEEE Aerospace Conference*, pages 319–324, Snowmass at Aspen, CO, USA, 1997.
- [12] E. Gat. On three-layer architectures. In D. Kortenkamp, R. P. Bonnasso, and R. Murphy, editors, *Artificial Intelligence and Mobile Robots*. AAAI Press, 1997.
- [13] M. P. Georgeff and F. F. Ingrand. Decision-making in an embedded reasoning system. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 972–978, Detroit, Michigan, 1989. AAAI Press, Menlo Park, CA.
- [14] M. P. Georgeff and A. L. Lansky. Reactive reasoning and planning. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 677–682. MIT Press, 1987.
- [15] S. E. Kleene. *Object-Oriented Programming in Common Lisp: A Programmer's guide to CLOS*. Addison-Wesley, 1988.
- [16] D. McDermott. The 1998 AI planning systems competition. *The AI Magazine*, 21(2):35–55, 2000.
- [17] D. McDermott, M. Ghallab, A. Howe, C. Knoblock, A. Ram, M. Veloso, D. Weld, and D. Wilkins. The PDDL planning domain definition language (available at <http://www.cs.yale.edu/>). Technical report, Yale Computer Science Department, 1998.
- [18] L. Spector and J. Hendler. The use of supervenience in dynamic-world planning. In K. Hammond, editor, *Proceedings of The Second International Conference on Artificial Intelligence Planning Systems*, pages 158–163, 1994.
- [19] R. St. Amant. *A Mixed-Initiative Planning Approach to Exploratory Data Analysis*. PhD thesis, University of Massachusetts, Amherst, 1996. Also available as technical report CMPSCI-96-33.
- [20] G. L. Steele Jr. *Common Lisp: The Language*. Digital Press, second edition, 1990.
- [21] D. L. Westbrook, S. D. Anderson, D. M. Hart, and P. R. Cohen. CLIP: Common lisp instrumentation package. Technical Report 94-26, University of Massachusetts at Amherst, Computer Science Department, Amherst, MA 01003, 1994.
- [22] D. E. Wilkins, K. L. Myers, J. D. Lowrance, and L. P. Wesley. Planning and reacting in uncertain and dynamic environments. *Journal of Experimental and Theoretical AI*, 7(1):197–227, 1995.