

Executing Reactive, Model-based Programs through Graph-based Temporal Planning

Phil Kim, Brian C. Williams, and Mark Abramson

Massachusetts Institute of Technology, Rm. 37-381

77 Massachusetts Ave. Cambridge, MA 02139 USA

E-mail: {kim,williams}@mit.edu

Abstract

In the future, webs of unmanned air and space vehicles will act together to robustly perform elaborate missions in uncertain and sometimes hostile environments. To achieve this robustness we go beyond current embedded programming languages, introducing a *model-based programming language* that enables autonomous vehicles to select and adapt coordinated mission plans on the fly. First, we present a variant of the Reactive Model-based Programming Language (RMPL), that allows for the expression of complex concurrent activities, metric time constraints and multiple contingencies. Second, we introduce the Temporal Planning Network (TPN), a simple, compact encoding of all possible executions of an RMPL program, that supports a model of program interpretation as fast temporal planning. Finally, we introduce an RMPL interpreter, called Kirk, which uses graph search on TPNs to find temporally consistent executions of RMPL programs.

1 Model-based Programming

The recent spread of advanced processing to embedded systems has created vehicles that execute complex missions with increasing levels of autonomy, in space, on land and in the air. These vehicles must respond to uncertain and often unforgiving environments, both with a fast response time and with a high assurance of first time success. The future looks to the creation of *cooperative robotic networks*. For example, giant space telescopes are being deployed that are composed of satellites carrying the telescope's different optical components. These satellites act in concert to image planets around other stars, or unusual weather events on earth. A heterogeneous collection of vehicles, such as planes, helicopters and boats, might work in concert to perform a search and rescue during a hurricane or similar natural disaster.

The creation of robotic networks cannot be supported by the current programming practice alone. Recent mission failures in the air and in space have highlighted the difficulty of creating highly capable vehicles within realistic budget limits. Due to cost constraints, flight software teams have not had time to think through all the plausible situations that might arise, encode the appropriate responses within their software and then validate that software with high assurance. To break through this barrier we need to invent a new programming paradigm.

In this paper we advocate the creation of *embedded, model-based programming languages*. First, programmers should retain control for the overall success of a mission, by programming game plans and contingencies that in the programmer's experience will ensure a high degree of success. The programmer should be able to program these game plans using features of the best embedded programming languages available. For example, reactive synchronous languages [Halbwachs, 1993], like Esterel, Lustre and Signal, offer a rich set of constructs for interacting with the sensors and actuators, for creating complex behaviors involving concurrency and preemption, and for modularizing these behaviors using all the standard encapsulation mechanisms. Model-based programming extends this style of reactive language with a minimal set of constructs necessary to perform flexible mission coordination, while hiding its reasoning capabilities under the hood of the language's interpreter or compiler.

Second, we argue that model-based programming languages should focus on elevating the programmer's thinking, by automating the process of reasoning about low-level system interactions. Most recent air and space mission failures can be isolated to human errors in reasoning through low-level system interactions. On the other hand, this limited form of reasoning and book keeping is the hallmark of computational methods. The interpreter or compiler of a model-based program reasons through these interactions using composable models of the system being controlled. We are developing a language, called the *Reactive Model-Based Programming Language (RMPL)*, that supports four types of reasoning about system interactions: reasoning about contingencies, scheduling, inferring a system's hidden state and

controlling that state.

This paper develops RMPL in the context of contingencies and scheduling. First, we introduce a subset of RMPL that includes constructs from traditional reactive programming plus constructs for specifying contingencies and scheduling constraints. Second, we compile RMPL programs to *temporal plan networks (TPN)*, which compactly represent all possible threads of execution of an RMPL program, and all resource constraints and conflicts between concurrent activities. Third, we present *Kirk*, an online interpreter for RMPL that “looks” by using network search algorithms to find threads of execution through the TPN that are temporally consistent. The result is a partially ordered temporal plan. Kirk then “leaps” by executing the plan using plan execution methods developed for Remote Agent [Tsamardinos *et al.*, 1998]. Finally, we discuss Kirk’s application to a simulated search and rescue mission.

2 Example: Cooperative Search and Rescue



As part of a search and rescue mission, consider an activity called Enroute, in which a group of vehicles fly together from a rendezvous point to the target search area. In this activity, the group selects one of two paths for traveling to the target area, flies together along the path through a series of waypoints to the target position, and then transmits a message to the forward air controller to indicate their arrival, while waiting until the group receives authorization to engage the target.

The two paths available for travel to the target area are each only available for a predetermined window of time, which is important to consider when selecting one of these paths. In addition, the timing of the Enroute activity is bound by externally imposed requirements, for example, the search and rescue mission must complete in 25-30 minutes, with 20% to 30% of the time allotted to the Enroute activity.

Codifying the Enroute activity requires most standard features of embedded languages. There are both sequential and concurrent threads of activities, such as going to a series of way points, and sending a message to the forward air controller (FAC), while concurrently awaiting authorization. There are maintenance conditions and synchronizations. For example, the air corridor needs to be maintained safe during flight, and synchronization occurs with the FAC.

In addition to constructs found in traditional embedded languages, we need constructs for expressing timing requirements and alternative choices or con-

tingencies, in this example to use one of two corridors. These constructs are common to robotic execution languages [Firby, 1995; Gat, 1996]. However, they are only used reactively. RMPL must reason forward through the program’s execution, identifying a course of action that is consistent.

3 RMPL Constructs

To summarize, RMPL needs to include constructs for expressing concurrency, maintaining conditions, synchronization, metric constraints and contingencies. The relevant RMPL constructs are as follows. We use lower case letters, like *c*, to denote activities or conditions, and upper case letters, like *A* and *B*, to denote well-formed RMPL expressions:

a. Invokes primitive activity *a*, starting at the current time. This is the basic construct for initiating activities.

c. Asserts that condition *c* is true at the current time, where *c* is a literal. This is the basic construct for asserting conditions.

if *c* thennext *A*. Starts executing *A* if condition *c* is currently satisfied, where *c* is a literal. This is the basic construct for expressing conditional branches and asserting preconditions.

do *A* maintaining *c*. Executes *A*, and ensures throughout *A* that *c* occurs. This is the basic construct for introducing maintenance conditions and protections.

A, B. Concurrently executes *A* and *B*. It is the basic construct for forking processes.

A; B. Consecutively executes *A* and then *B*. It is the basic construct for sequential processes.

A[*l, u*]. Constrains the duration of program *A* to be at least *l* and at most *u*. This is the basic construct for expressing timing requirements.

choose {*A, B*}. Reduces non-deterministically to program *A* or *B*. This is the basic construct for expressing multiple strategies and contingencies.

Note that together, *c* and **if *c* thennext *A*** provide the basic constructs for synchronization, by specifying required and asserted conditions. *A, B* and *A; B* provide the necessary constructs for building complex concurrent threads. Using these constructs we express the Enroute activity as follows:


```
Group-Enroute() [1,u] = {
  choose {
    do {
      Group-Fly-Path(PATH1_1,PATH1_2,
        PATH1_3,TAI_POS) [1*90%,u*90%];
    } maintaining PATH1_OK,
    do {
      Group-Fly-Path(PATH2_1,PATH2_2,
        PATH2_3,TAI_POS) [1*90%,u*90%];
    } maintaining PATH2_OK
  };
  {
    Group-Transmit(FAC,ARRIVED_TAI) [0,2],
    do {
      Group-Wait(TAI_HOLD1,TAI_HOLD2)
        [0,u*10%]
    } watching PROCEED_OK
  }
}
```

The *choose* expression models the two options for flight paths. 90% of the total time of the overall maneuver is allocated to this group flight. Each flight has a maintenance condition that the flight path is okay. Arrival is transmitted to the forward air controller, and receipt of a message to proceed is concurrently monitored.

4 Temporal Plan Networks

Executing an RMPL program involves choosing a set of threads of execution (*Plans*), checking to ensure that the execution is consistent and schedulable, and then scheduling events on the fly. It is essential that we generate these plans quickly. This suggests compiling RMPL programs to a plan graph, along the lines of Graphplan or Satplan,[Weld, 1999], and then search the pre-compiled graph. However, it is also important for the plan to have the temporal flexibility offered by a partially ordered, temporal plan. Least commitment leaves slack to adapt to execution uncertainties and to recover from faults. This partial commitment is expressed in temporal planning through a *Simple Temporal Network* (STN)[Dechter *et al.*,]. Hence, a key observation of our approach is that to build in temporal flexibility we should build our graph-based plan representation, called a *Temporal Plan Network* (TPN), as a generalization of an STN.

The TPN corresponding to the above Enroute program is shown below. Activity name labels are omitted to keep the figure clear, but the node pairs 4,5 and 6,7 represent the two Group-Fly-Path activities, and node pairs 9,10 and 11,12 correspond to the Group-Wait and Group-Transmit activities, respectively. Node 3 is a decision node that represents a choice between two methods for flying to the search area. The TPN represents the consequences of the constraint that the mission last between 25 and 30 minutes. It also models the decision between the two paths to the target area, and it models the restrictions that each of the paths can only be used if they are available.



A TPN encodes all feasible executions of an activity. It does this by augmenting an STN with two types of constraints: temporal constraints restrict the behavior of an activity by bounding the duration of an activity, time between activities, or more generally the temporal distance between two events. Symbolic constraints restrict the behavior of an activity by expressing the assertion or requirement of certain conditions by activities that all valid executions must satisfy.

For example, consider some of the possible executions

of the Enroute activity. One possible execution is that the group flies along path one to the target area in 420 time units (seconds in this case), transmits an arrival message to the forward air controller in one second, then waits for another 40 seconds to receive authorization to proceed. Another possible execution is that the group selects the second path, flies to the target area in 500 seconds, takes 2 seconds to transmit the arrival message, and is authorized to proceed immediately. If it were the case that path one was available from the time at which the Enroute activity started to at least the time that the group arrived at the target area, then the first execution is valid. This is because it satisfies both the temporal constraints on the Enroute activity, and the requirement that path one is available for the duration of the flight along it. The planning algorithm presented in the next section performs the identification of consistent activity executions.

A Temporal Planning Network is a Simple Temporal Network, augmented with *symbolic constraints* and *decision nodes*. These additions are sufficient to capture all RMPL constructs given earlier. Like a simple temporal network, the nodes of a TPN represent temporal events, and the arcs represent temporal relations that constrain the temporal distance between events. An arc of a TPN may be labeled with a symbolic constraint Tell(c) or Ask(c), as well as a duration. A Tell(c) label on an arc (i,j) asserts that the condition represented by c is true over the interval between the temporal events modeled by the nodes i and j. Similarly, an Ask(c) label on an arc (i,j) requires that the condition represented by c is true over the interval represented by this arc. For example, in the Enroute TPN, the Ask(PATH1=OK) label on the arc (3,4) represents the requirement for path one to be available for the interval of time corresponding to the interval of time between the temporal event modeled by node 3 and node 4. These Ask-type symbolic constraints allow for the encoding of conditions in the network.



Decision nodes are used to explicitly introduce choices in activity execution that the planner must make. For example, in the Enroute activity there are two choices of paths for the group to use for flying to the target area, path one and path two. The activity model captures the two choices as out-arcs of node 3 of the enroute TPN. This decision node is designated by a double outline and dashed out-arcs. All other nodes in the Enroute TPN are non-decision nodes.

The addition of Tell and Ask constraints and decision nodes offer a simple but powerful extension to STNs that is sufficient to capture the rich set of RMPL constructs listed earlier.

5 Compiling RMPL to TPN

Given a well formed RMPL expression, we compile it to a TPN by mapping each RMPL primitive to a TPN as defined below. RMPL sub-expressions, denoted by upper case letters, are recursively mapped to equivalent

$A[l, u]$. Invoke activity A between l and u time units.



Given an RMPL program, the Kirk planner uses its compiled TPN to search for an execution that is both complete and consistent. The execution corresponds to an unconditional, temporal plan. A plan is complete if choices have been made for each relevant decision point, it contains only primitive-level activities, and all activities labeled Ask(c) have been linked to a Tell(c). A

Network search completes only when all paths reach the end-node of the top-level activity, and the subnetwork of the TPN, defined by these paths, is temporally

consistent. This corresponds to testing consistency of an STN[Dechter *et al.*,], as discussed in the next section.

The first phase of planning is summarized by the *Modified Network Search algorithm*, shown below. The set A, is the set of active nodes, which are those nodes whose paths have not yet been fully extended. The sets SN and SA are the sets of selected nodes and selected arcs, respectively:

```

1 Modified-Network-Search( N )
2   A = { start-node of N };
3   SN = { start-node of N };
4   SA = { };
5   While ( A is not empty )
6     Node = Select and remove a member of A;
7     If ( Node is a decision-node )
8       Arc = Select any unmarked out-arc of Node and
9       Mark Arc and
10      Add Arc to SA;
11      If ( tail of Arc is not in SN )
12        Add tail of Arc to A and SN;
13      End-If
14    Else
15      For each Arc that is an out-arc of Node
16        Add Arc to SA;
17        If ( tail of Arc is not in SN )
18          Add tail of Arc to A and SN;
19        End-If
20      End-For
21    End-If
22  If ( Cycle-Induced(SN, SA) )
23    If ( Not(Temporally-Consistent(SN, SA)) )
24      Backtrack(SN, SA, A);
25    End-If
26  End-While
27 End-Function

```

The algorithm extends an active node at each iteration. Decision nodes are treated by extending the path along one out arc (lines 8-13), while non-decision nodes are treated by branching the path and extending along all out arcs (lines 15-20). At the end of each iteration of the main While-loop, the modified network search tests for temporal consistency (lines 24-26). If the test fails, then the search calls Backtrack(..) in line 25, which reverts SN, SA, and A to their states before the most recent decision that has unmarked choices remaining, and selects a different out-arc. While for simplicity this explanation uses chronological backtracking, a wealth of more efficient search algorithms can be applied.

Note that it is not necessary to check temporal consistency after every iteration of the While-loop, since as long as no cycles are induced in the network, there is no way for a temporal inconsistency to be induced. Determining whether a cycle has been created can be done for each arc that is selected by checking whether the arc's tail node has already been selected. Since this can be done in constant time, it is significantly more efficient in practice than testing temporal consistency after every iteration, although it doesn't impact worst case complexity.


Also note that the algorithm stops extending a path when it encounters a node that is already in SN. The fact that this node is already in SN implies that two concurrent threads of execution have merged.

Finally, after the modified network search completes,


the selected nodes and arcs define a set of paths from the start-node to the end-node of the top activity.

Example: Searching the Enroute Network


To illustrate the modified network search, we return to the Enroute input network, where node 1 is the start-node and node 2 is the end-node:




Initially, node 1 is selected, which is indicated by its darker shade, and it is active, which is indicated by its dashed outline. In the first iteration, Kirk chooses node 1 from the set of active nodes, and since node 1 is not a decision node, it selects all out-arcs and adds their tails to the selected and active set. This continues until both node 5 and node 15 are selected:



At this point, the modified network search chooses node 5 from the active set. Since node 5 is a decision node, the algorithm must choose either arc (5,7) or arc (5,10). It selects arc (5,7) and continues extending until it reaches the following:



Note that arc (14,2) is selected, forming the cycle, 1-3-4-5-7-8-9-6-13-14-2-1, so the algorithm checks for temporal consistency. In this example, this selected sub-network is temporally inconsistent, so the algorithm backtracks to the most recent decision with open options, which is Node 5. Out-arc (5,10) has not yet been tried, so it is selected and the path extend to the end-node. Finally a path through arc (15,16) is found to the end-node, resulting in the temporally consistent sub-network:



Checking Temporal Consistency

To check temporal consistency we note that any subnet of a Plan Network, minus its symbolic constraint labels, forms a Simple Temporal Network. Hence temporal consistency can be checked using standard methods for

Simple Temporal Networks [Dechter *et al.*,]. Recall that an STN is consistent if and only if its encoding as a distance graph contains no negative cycles [Dechter *et al.*,]. There exist several well known algorithms for detecting negative cycles in polynomial time. The Bellman-Ford algorithm [Cormen *et al.*, 1990] can be used to check for negative cycle in $O(nm)$ time, where m is the number of arcs in the distance graph. This algorithm only needs to maintain one distance label at each node, which takes only $O(n)$ space. A variant of this algorithm is used by HSTS [Muscettola *et al.*, 1998] for fast inconsistency detection.

The algorithm we use in the Kirk planner is a variant of the generic label-correcting single-source shortest-path algorithm [Ahuja *et al.*, 1993], which takes $O(nm)$ worst-case asymptotic running time, but performs faster in many situations. This algorithm also requires only $O(n)$ space. Space precludes a more detailed development.

6.2 Phase Two: Threats and Open Conditions

Symbolic constraints— Ask(c) and Tell(c) — are handled analogous to threats and open conditions in causal link planning[Weld, 1994]. Two symbolic constraints conflict if one is either asserting (by using Tell) or requesting (by using Ask) that a condition is true, and the second is asserting or requesting that the same condition is false. For example, Tell(Not(c)) and Ask(c) conflict. An open condition in a TPN appears as Ask constraints, which represent the need for some condition to be true over the interval of time represented by the arc labeled with the Ask constraint.

Resolving Threats


To detect threats the planner computes the feasible time bounds for each temporal event (node) in the network, and then uses these bounds to identify potentially overlapping intervals that are labeled with inconsistent constraints. These bounds can be computed by solving an all-pairs shortest-path problem over the distance graph representation of the partially completed plan [Dechter *et al.*,]. Kirk used the Floyd-Warshall algorithm for computing all-pairs shortest paths because of ease of implementation. We are currently evaluating Johnson's algorithm which runs in $O(n^2 \log(n) + mn)$, or $O(n^2 \log(n))$ if $m = O(n)$.

Once these feasible time ranges are determined, the planner detects which arcs may overlap in time. If there are two arcs that may overlap and that are labeled with conflicting symbolic constraints, then they are resolved by ordering the intervals, if possible.

These interval pairs need to be identified efficiently. Kirk maintains an interval set data structure for each proposition p that keeps track of all intervals that assert or require p or its negation. In order to identify threats, the planner need only check each interval set for threats. This takes $O(si^2)$ asymptotic running time, where i is the maximum cardinality over all interval sets, and per-


forms much better in practice because the interval sets typically have few elements. More sophisticated indexing schemes may improve performance, such as interval trees structures [Cormen *et al.*, 1990].

A threat is resolved by introducing temporal constraints. Each threat consists of two arcs that represent intervals of time that may overlap. To resolve threats we introduce a constraint that forces an ordering between the two activities, similar to promotion and demotion in classical planning[Weld, 1994]:



Closing Open Conditions

An open condition is represented by an arc labeled with an Ask constraint, which represents the request for a condition to be satisfied over the interval of time represented by the arc. If this interval of time is contained by another interval over which the condition is asserted by a Tell constraint, then the open condition is satisfied (i.e., closed), and a causal link is drawn from the Tell to the Ask. Open conditions are detected simply by scanning through all activities and checking any Ask constraints. Finding potentially overlapping intervals is performed using the same method described above for detecting threats. Once a Tell is found that can satisfy an open condition, temporal constraints are added so that the duration of the open condition is contained within the Tell. This method of closing open asks is also closely related to the way that the HSTS planner satisfies compatibilities [Muscettola *et al.*, 1998]:



7 Implementation and Results

The RMPL Compiler generates TPN specification files, and is written in Lisp. Kirk, written in C++, generates a plan from the TPN and checks consistency. The Plan Runner [Tsamardinos *et al.*, 1998], takes the resulting partially ordered temporal plan and executes it on the multi-air vehicle simulator. The results is visualized using world tool kit (as shown earlier). The following table summarizes Kirk's performance on nominal plans for several activities within the search and rescue scenario. The fully expanded TPN generated from the Group-Search-and-Rescue activity included 273 nodes. The testing platform was an IBM Aptiva E6U with an

Intel 400Mhz Pentium II processor and 128MB of RAM, running Redhat Linux version 6.1:

Top Activity	Nodes	Activities	Plan Time
Follow(..)	4	1	4 ms
Group-Rescue(..)	27	8	235 ms
Group-Enroute()	112	19	16 s
Group-SR-Mission()	273	47	404 s

“Top Activity” refers to the top-level activity that was being planned. “Nodes” is the size of the expanded TPN after planning. Usually, about half of these were included in the final plan, with the rest corresponding to unselected executions. “Activities” indicates the number of primitive activities included in the final plan. Finally, the “Plan Time” gives the time that it took for Kirk to generate a plan corresponding to each of these activities.

Temporal planners, such as the DS1 remote agent planner, can take many hours to generate a plan of similar size, unless hand coded search heuristics are provided. Kirk does well with no search guidance up to about 100 nodes. At this point the time becomes dominated by the time required to compute feasible time bounds for events. This is due to the use of Bellman-Ford and chronological search in the first prototype, summarized above. Our reimplementation is based on Johnson’s algorithm and a more sophisticated search strategy. Preliminary results suggest a significant performance increase. We are currently extending Kirk to support both decision theoretic planning and agile path planning.

Acknowledgments

This research is supported in part by the Office of Naval Research under contract N00014-99-1-1080.

References

- [Ahuja *et al.*, 1993] R. Ahuja, T. Magnanti, and J. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [Cormen *et al.*, 1990] T. Cormen, C. Leiserson, and R. Rivest. *Introduction to Algorithms*. MIT Press, Camb., MA, 1990.
- [Dechter *et al.*,] R. Dechter, I. Meiri, and J. Pearl. Temporal constraint networks.
- [Firby, 1995] R. James Firby. The RAP language manual. Technical report, Univ. Chicago, 1995.
- [Gat, 1996] Erann Gat. Esl: A language for supporting robust plan execution in embedded autonomous agents. In *AAAI Fall Symposium on Plan Execution*, 1996.
- [Halbwachs, 1993] N. Halbwachs. *Synchronous programming of reactive systems*. Kluwer Academic, 1993.
- [Muscettola *et al.*, 1998] N. Muscettola, P. Morris, B. Pell, and B. Smith. Issues in temporal reasoning for autonomous control systems. In *Autonomous Agents*, 1998.
- [Tsamardinos *et al.*, 1998] I. Tsamardinos, N. Muscettola, and P. Morris. Fast transformation of temporal plans for efficient execution. In *Proceedings of AAAI-98*, 1998.
- [Weld, 1994] D. Weld. An introduction to least commitment planning. In *AI Magazine*, 1994.
- [Weld, 1999] D. Weld. Recent advances in ai planning. In *AI Magazine*, 1999.