

1

AD-A256 198



A Specification of the Soar Cognitive Architecture in Z

Brian G. Milnes

with contributions from

Garrett Pelton, Robert Doorenbos, Mike Hucka,
John Laird, Paul Rosenbloom and Allen Newell

August 31, 1992

CMU-CS-92-169

DTIC
ELECTE
OCT 07 1992
S A D

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

B. Milnes, G. Pelton, B. Doorenbos may be reached at the School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213-3890. A. Newell is deceased, but his academic estate may also be reached at the School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213-3890. M. Hucka and J. Laird may be reached at the Electrical Engineering and Computer Science Department, The University of Michigan, Ann Arbor, MI 48109-2122. P. Rosenbloom may be reached at the Information Sciences Institute and Department of Computer Science, University of Southern California, 4676 Admiralty Way, Marina del Rey, CA 90292.

This document has been approved for public release and sale; its distribution is unlimited.

This research was sponsored by the Avionics Laboratory, Wright Research and Development Center, Aeronautical Systems Division (AFSC), U.S. Air Force, Wright-Patterson AFB, Ohio 45433-6543 under Contract F33615-90-C-1465, ARPA Order No. 7597.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of DARPA or the U.S. government.

92

423887

92-26550



24308

Keywords: Artificial Intelligence, formal specification, integrated intelligent architectures, machine learning, problem solving, Z specifications, Z notation, Soar.

Abstract

A formal specification of the sixth revision of the Soar architecture in the Z notation was constructed to elucidate and clarify the definition of Soar and to guide its implementation. Soar is a cognitive architecture that has been successfully applied to many domains and has been proposed as an exemplar unified theory of cognition. Z is a model theoretic specification language based in set theory that has syntax and type checking programs available. The specification has a complete coverage of the architecture, a low level of abstraction and a considerable implementation bias.

Accession For	
NTIS CR&I	
DTIC TAB	
Unannounced	
Justification	
By	
Distribution	
Availability	
Dist	Avail
A-1	Soar

ED 1

For Allen Newell (1927-1992) — who taught me that computer science is an empirical quantitative discipline.

—B.G.M.

Contents

1	Introduction	10
1.1	An Introduction and the Intended Audience	10
1.2	The History and Philosophy of the Specification	11
1.3	The Specification's Coverage and Implementation Bias	12
1.4	An Overview of the Architecture	12
1.5	An Overview of the Specification	16
2	Specifying Soar's Control: State Machines	18
2.1	Specifying Soar's Control	18
2.2	The Prototype State Machine	19
3	Base Symbol Structures	30
4	Temporary Memory and Production Memory	33
4.1	Preferences	33
4.2	Working Memory Elements	37
4.3	Temporary Memory Elements	39
4.4	Makes	40
4.5	Tests	43
4.6	WME Tests	46
4.7	Conditions	47
4.8	SP	50
5	Recognition Memory	54
5.1	Bindings	55
5.2	Matching	55
5.2.1	Tests Matching Symbols	56
5.2.2	Matches	58
5.2.3	Matching against Working Memory	59
5.3	Instantiation	60
5.3.1	Instantiating a Match's Production's Conditions	60
5.3.2	Instantiating the RHS of a Match's Production	66

5.3.3	Instantiation	67
5.4	Goal Operations and a Match's Goal	68
5.5	Chunking	69
5.5.1	Extracting Not Tests	70
5.5.2	Variabilize	72
5.5.3	Not-ify	75
5.5.4	Traces	79
5.5.5	Chunking's State	81
5.5.6	Changing, Initializing and Resetting Chunking	86
5.5.7	Defining Results	88
5.5.8	Chunking an Instantiation	88
5.5.9	Tracing Seeds, Locals and Grounded Potentials	91
5.5.10	Tracing UnGrounded Potentials	97
5.5.11	ChunkSP	100
5.5.12	Step Chunking	101
5.6	Slots	101
5.7	Recognition Memory's State	102
5.8	Support	106
5.9	Preference Phase Operations	110
5.9.1	Initialize, Finish, Reset and Quiescence	111
5.9.2	Retracting Instantiations	111
5.9.3	Instantiate Match	112
5.9.4	Conferring Support	114
5.9.5	Stepping Chunking	114
5.9.6	O-rejects	115
5.9.7	Changing Preference Memory	116
6	IO	119
6.1	Channels, Attributes and IO Mappings	119
6.2	Output Structures	121
6.3	Input Structures	121
6.4	Cycles of IO	122
6.5	IO State	122
6.6	The Input Cycle	125
6.6.1	Legal Input	126
6.6.2	Initializing and Starting the <i>InputCycle</i>	128
6.6.3	Reading <i>NewInputStructures</i>	128
6.6.4	Reading <i>ModifyInputStructures</i>	130
6.6.5	Reading <i>DeleteInputStructures</i>	132
6.6.6	Reading an Input Channel	132
6.6.7	Finishing, Stepping and Resetting the <i>InputCycle</i>	133
6.7	The Output Cycle	134

7	Decide	137
7.1	Impasses	138
7.2	Impasser State	139
7.3	Preference Semantics State	142
7.4	Working Memory Phase State	144
7.5	Quiescence Phase State	146
7.6	Decide's Total State	149
7.7	Impasser Transitions	151
7.8	Preference Semantics Transitions	156
7.9	Working Memory Phase Transitions	167
7.10	Quiescence Phase Transitions	173
8	Top Level	185
8.1	The States of the Top Level	185
8.2	Soar's State Schema	186
8.3	The Top Level State Machine	188
A	Numbers	194
B	An Implementation Discipline	196
B.1	Observing Soar's Operation	196
B.2	An Implementation Discipline	197
B.3	An Estimate of the Cost of Using State Machines	198
C	Finer Grained Hooks	201
C.1	TM Hooks	201
C.1.1	WME Operations	202
C.1.2	Preference Operations	202
C.1.3	TME Operations	203
C.2	SPM Hooks	203
C.3	GM and IM Hooks	204
C.3.1	Context Impasses	204
C.3.2	Non-Context Impasses	204
D	The Reorderer	206
D.1	Classifying WMETests	206
D.2	Deriving the Branching Factor Heuristic	209
D.3	The Reorderer's State	218
D.4	The Reorderer's Operations	221

List of Figures

1.1	Module and Control Block Diagram	13
1.2	Module and Data Block Diagram	14
2.1	State Machine Prototype	20
5.1	Chunking's State Machine	82
5.2	Preference Phase State Machine	104
6.1	Graphical Representation of IO	120
6.2	Input Cycle State Machine	123
6.3	Output Cycle State Machine	124
7.1	Impasser State Machine	140
7.2	Preference Semantics State Machine	143
7.3	Working Memory Phase State Machine	145
7.4	Quiescence Phase State Machine	147
8.1	Top Level State Machine	187
D.1	Reorderer State Machine	219

Chapter 1

Introduction

1.1 An Introduction and the Intended Audience

A formal specification of the sixth revision of the Soar cognitive architecture [LNR87] was constructed in the Z specification notation [Spi89, Spi88]. Soar is a cognitive architecture that has exhibited a wide variety of performance in domains such as: machine learning [LRN84], cognitive modeling [Aky90, DS90, JRS91], development [SK91], natural language processing [LLN91], expert systems [HPS89] and robotics [LYHT91].

This specification was written to guide the construction of implementations of Soar and to facilitate the understanding and extension of the Soar architecture. Formal specification has simultaneously been applied to two other related computational architectures by Iain Craig of the University of Warwick, UK [Cra91]. Craig used Z to specify blackboard architectures and his Casandra architecture and similarly applied the specification to guide Casandra's implementation.

The Z notation is introduced as it is used. It is largely founded upon familiar set theory, but most readers will be better served by first reading an introductory text, such as [PTS91, Wor92, Hay85], and the Z manual [Spi89].

Artificial intelligence researchers who have an interest in understanding the details of Soar will find many answers to their questions here. Indeed, the process of writing the specification has made clear many details of the architecture to its implementors and founders. Considerable knowledge of the Soar architecture from less formal sources, such as [LCAS90], would ease and enhance the reader's understanding. Z specifiers will most likely find this difficult reading, but we do draw some conclusions about the strengths and weaknesses of using Z for a large, detailed, low-level specification of a computational architecture.

1.2 The History and Philosophy of the Specification

In January 1990 we began an effort to re-implement the current Soar architecture, intending only a few well understood modifications. As we proceeded, it quickly became clear that the limiting factor was our understanding of the details of the architecture itself. Although Soar's architecture description papers [LNR87, LRN86] and manual [LCA590] described some of the architecture, they were out of date with the current system (Soar 5.2), at too high a level of abstraction and had obvious inconsistencies with the system.

In an earlier revision of Soar (4.2) a small module (preference semantics) had been formalized in first order logic [Lai86]. When we updated and reorganized this formalization, we discovered that it had a pleasing clarity and disambiguity that allowed us to quickly answer precise questions about the module's operation.

After considerable reflection, we decided to undertake a formal specification of the basic functionality of the architecture. Our hope is that this document will become the definitional standard of Soar. That is, the specification's terms will become the language that we use to discuss the architecture. And, when we change the architecture, we will first update the specification to flush out our understanding of the changes, and then distribute the changed specification for comment. This way, the detailed content of the architecture will not remain hidden in the implementation and become obscured with time.

After an initial, unsuccessful attempt to formally specify the architecture in an ad hoc language, we undertook a small survey of the popular specification languages. Our intuition was that a model theoretic language would serve the purpose of specifying a cognitive architecture better than an algebraic system such as Larch [Gut90]. We looked at Z, VDM [Jon90, JS90], and OBJ [GW88], but chose Z. Z appealed most to us because of its strengths in composing small specifications into larger specifications, the availability of good programs to textually format, syntax and type check its specifications and the local expertise here at Carnegie Mellon with Z [Gar91, DG89].

Although Z has proved a very capable specification language it lacks two properties that would have considerably simplified our task. First, Z's lack of a full execution semantics forced us to use a cumbersome state machine notation. Second, although we agree that not all specifications written in a language should be executable, if Z was directly interpretable it would have allowed us to much more quickly debug the specification.

1.3 The Specification's Coverage and Implementation Bias

The specification is intended to cover all the functionality of what we view as "the Soar architecture", and avoid details that are specific to the implementation. Initially we felt that the ideal specification would be a very abstract specification of the architecture with little implementation bias because an abstract specification would simplify the modification of the specification and help to prevent discourse on the design from being overwhelmed in details.

Several problems quickly derailed our aspirations for an abstract specification. We wanted the specification to guide our re-implementation of Soar. We desired an easily instrumented and flexibly modified implementation. This drove us to produce a more detailed specification that could be used as an index into the implementation. Also, as we have never had any specification of the architecture that was more concrete than general prose but less concrete than code, we found it difficult to produce abstract specifications that had sufficient detail to guide the implementation.

The result is a specification that abstracts away from many implementation details but is a very detailed accounting of the operation of the system. Some implementation oriented details, such as techniques for inserting "hooks" into the system, Soar's use of numbers, and a conjunctive match condition ordering algorithm have been included in appendixes. We leave the problem of how to produce a more abstract specification to future Soar researchers as they attempt to understand, modify and implement new revisions of the architecture.

1.4 An Overview of the Architecture

Figure 1.1 diagrams a module and control oriented view of Soar, and Figure 1.2 diagrams a module and data oriented view. Soar is composed of four modules and is enclosed in an interface, which we call the external interface. Figure 1.1 represents each module by a rectangle, each state machine control by a circle or oval, and the calling hierarchy by arrows. Figure 1.2 represents each module in a rectangle, each state machine with a circle or oval, each memory with a heavy lined rectangle, and the reading and writing of each memory with arrows.

The External Interface Soar is designed as a situated autonomous agent that directly interacts with the external world, but an actual implementation of the architecture is a program that allows a user to run Soar as if it is an embedded agent. We distinguish Soar the program, from Soar the architecture, by wrapping Soar in the external interface. This interface contains the user interface that allow users to view, control and modify the execution of Soar. The architecture provides only one operation to the external interface, *StepTL*.

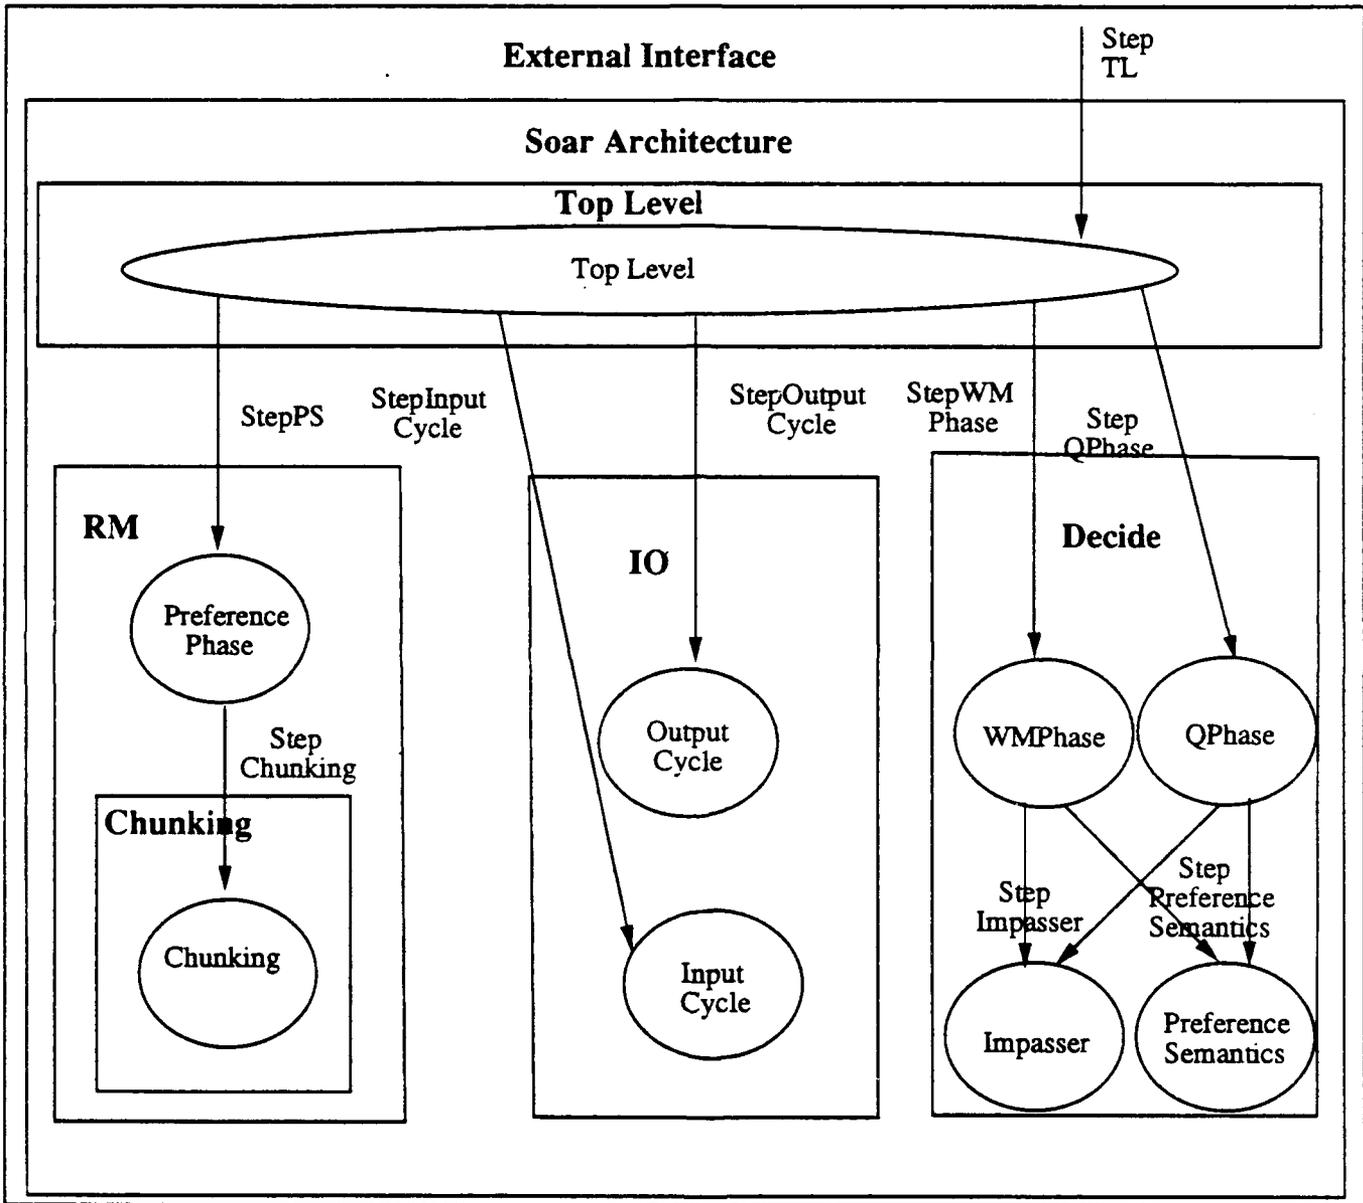


Figure 1.1: Module and Control Block Diagram

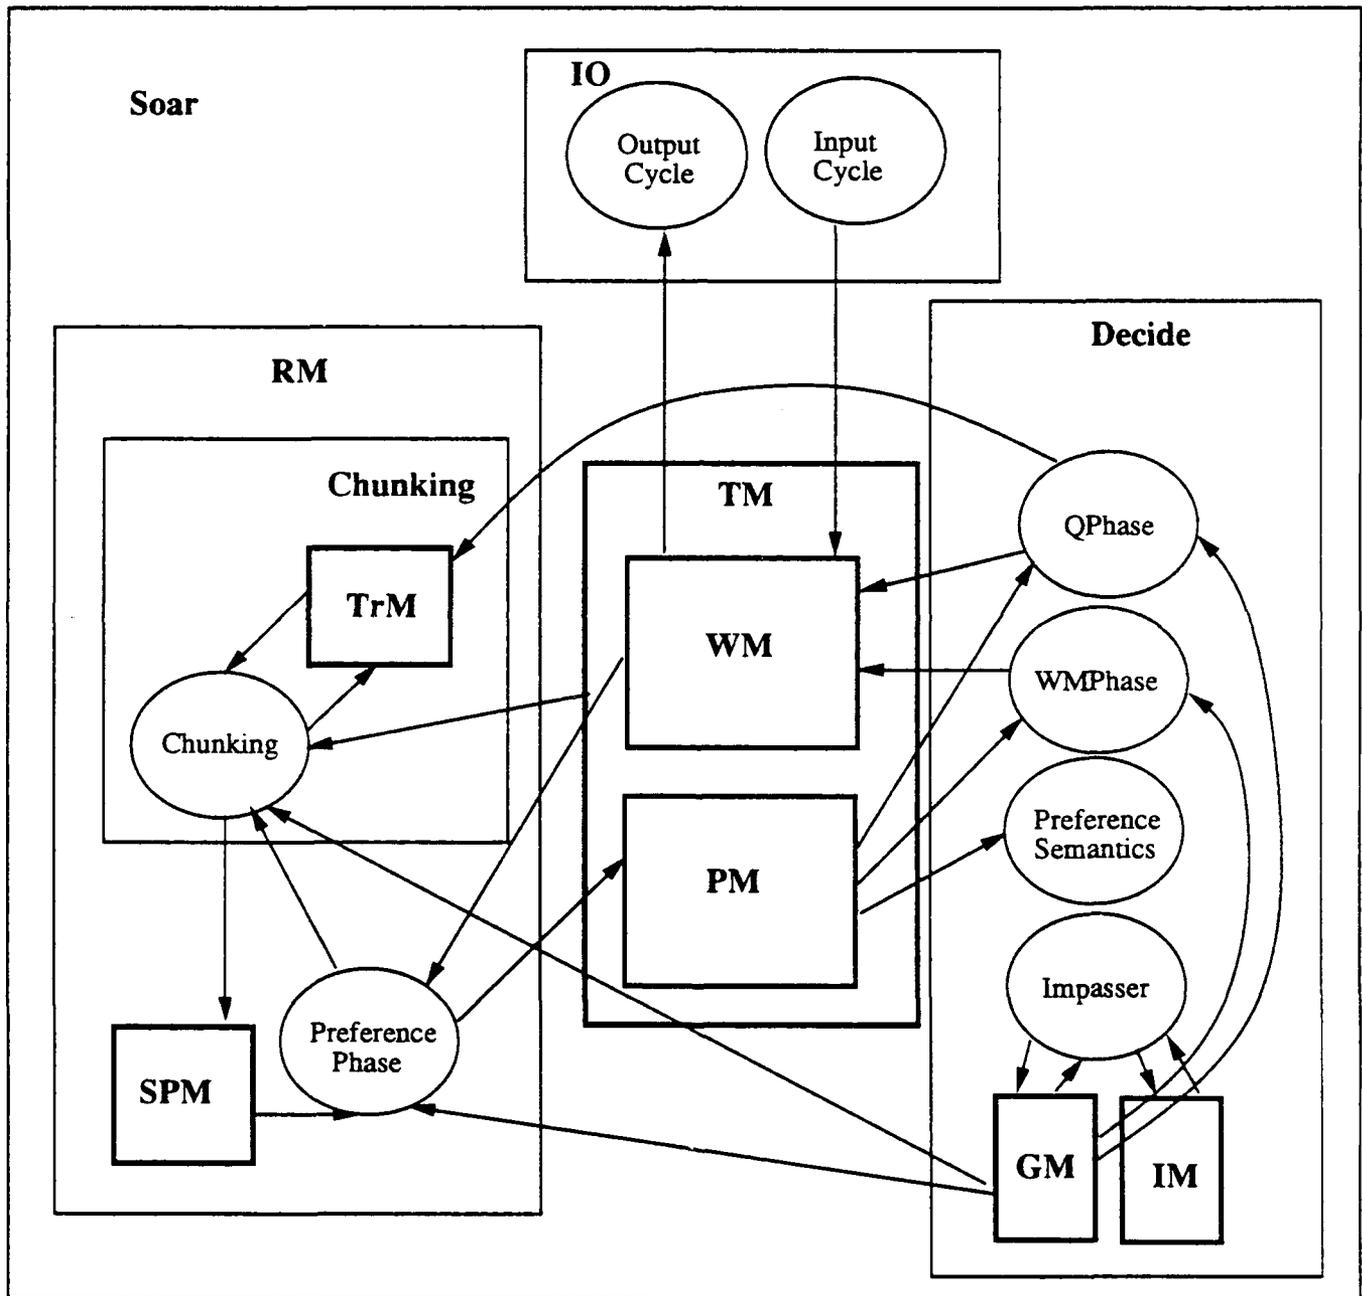


Figure 1.2: Module and Data Block Diagram

StepTL moves Soar one step forward. It acts much like a clock line of a CPU that an engineer externally pulses.

Soar's four modules are: the top level, recognition memory (RM) including chunking, input output (IO) and the decision procedure (Decide).

Recognition Memory Recognition memory is Soar's memory model; it contains a long term memory in the form of production rules called production memory, or *SPM*. Its short term memory is temporary memory, *TM*, which is split into two parts called working memory, *WM*, and preference memory, *PM*. As all of the modules of Soar read and write temporary memory, Figure 1.2 depicts it as global to the entire system. The preference phase is a loop that recalls information from long term memory until no new recollections are available.

Chunking Chunking is Soar's learning mechanism. It observes the changes to temporary memory and preference phase's recollections and stores them in two trace memories (*TrM*). From these and Decide's goal stack, it learns new memories and adds them to production memory. The external interface steps the top level, which calls *StepPS* to step recognition memory's preference phase, which calls *StepChunking* to step the chunking module.

IO The IO module allows Soar to communicate with the outside world. The top level steps IO's input and output cycles with *StepInputCycle* and *StepOutputCycle*. The input cycle reads perceptions from input channels, and adds them to temporary memory. The output cycle finds the subsets of temporary memory that represent motor actions and ships them out output channels to the outside world.

IO requires transducers to map Soar's motor commands into a form that simulated or real motor controllers can understand, and to map the perceptions of simulated or real perceptual devices into Soar's form. These transducers are housed in the external interface, and are also viewed as being part of SoarIO, but not part of the architecture.

Decide The decision procedure is Soar's universal subgoal mechanism. It maintains two goal memories, *GM* and *IM*, and interprets the contents of preference memory as requests to change the contents of pieces of working memory. Whenever a set of requested changes are not consistent, or do not provide a unique change to working memory, decide generates a new goal to represent the problem. The decision procedure has two state machines: working memory phase (*WMPhase*) and quiescence phase (*QPhase*), stepped by *StepWMPhase* and *StepQPhase*. The working memory phase and the quiescence phase both step the impasser and preference semantics, with *StepImpasser* and *StepPreferencePhase*.

1.5 An Overview of the Specification

The specification is organized into eight chapters and is supplemented with four appendixes on implementation related considerations.

1. Chapter 1, Introduction — this introduction and overview of the specification.
2. Chapter 2, Specifying Soar's Control: State Machines — introduces the state machine convention for specifying the high level control of Soar.
3. Chapter 3, Base Symbol Structures — introduces the basic symbol types which Soar composes into more complicated symbol structures.
4. Chapter 4, Temporary Memory and Production Memory — defines the temporary memory and production memory data structures that are used throughout the specification.
5. Chapter 5, Recognition Memory — defines the operation of the recognition memory, including Soar's learning mechanism, Chunking.
6. Chapter 6, IO — specifies Soar's mechanism for communicating with the external world.
7. Chapter 7, Decide — defines Soar's universal subgoaling mechanism.
8. Chapter 8, Top Level — specifies the top level control loop that orchestrates the execution of Soar's modules.
9. Appendixes —
 - (a) Appendix A, Numbers — specifies how to augment the base symbol structures of Soar with a single type of numeric type.
 - (b) Appendix B, An Implementation Discipline — describes techniques for creating an implementation that is very close to the specification.
 - (c) Appendix C, Finer Grained Hooks — describes a technique that would allow an implementation to observe and modify the implementation at a finer granularity than would otherwise be supported by the specification.
 - (d) Appendix D, The Reorderer — specifies a greedy heuristic reordering algorithm to sort the conditions of individual Soar productions into an order that on average matches faster.

Chapters 2, 3 and 4 cover introductory structure and are required to understand the remaining chapters. Chapters 5, 6, 7 and 8 cover the details of the major components of Soar, and may be read in any order. Although the Z

specification language requires a bottom-up definitional approach. chapters 5, 7 and 8 refer forward in prose and backwards in *Z* and prose. Two readings of these chapters is probably required to clearly understand their complicated interactions.

Chapter 2

Specifying Soar's Control: State Machines

This chapter describes how the structure of Z and the desired properties of the design and implementation of Soar conspire to produce many constraints on the specification of Soar. Section 2.1 discusses the alternative representations of control in Z. Section 2.2 describes the technique adopted to handle this problem: state machines.

2.1 Specifying Soar's Control

Z's technique for modeling a computational system is called a "sequential system model" ([Spi89] 130-141). Unfortunately, Z's sequential system model does not provide sufficiently rich execution semantics. As Soar is a computational model of cognition, a large part of its design is its flow of control. Although a specification that abstracts away from Soar's detailed control could prove useful, to guide our implementation effort, we would prefer that our specification carefully specify Soar's flow of control.

We have investigated three alternative techniques for specifying Soar's control in Z: prose, execution sequences and state machines.

We began by specifying pieces of the architecture's flow of control in prose. This prose described how individual state changes were combined in sequences and loops. However, it lacked clarity, coverage and precision.

Next, we experimented with specifying control by defining execution sequences. We constructed the set of all sequences of states of Soar, and then constrained them to start with an initial state and only follow defined state changes. However, as Soar contains several components each of which contains isolated control, this quickly became cumbersome. Finally, we settled upon a state machine control paradigm.

The state machine control makes explicit the legal sequencing of operations in the Soar model by using state diagrams, a Z representation of the states of the machine and a Z sequential operation for each transition of the state machine. The individual transitions of the state machine are composed into a single operation that moves the machine forward one of the possible transitions with each application.

In this chapter we construct a prototype state machine format that allows state machines to:

- initialize
- take one step of a possible set of transitions
- step an embedded sub-state machine
- determine when a state machine has finished executing and
- reset

The individual state machines that control parts of Soar are constructed following the form of this prototype, and then composed into a single machine that drives all of Soar. The resulting implementation of the architecture is controlled from an interface that we call the external interface. This interface drives Soar step by step, much like an engineer externally pulses the clock of a computer processor.

2.2 The Prototype State Machine

This section defines the state machine control conventions by exhibiting a state machine prototype. The prototype state machine is diagrammed in Figure 2.1. The states of the machine are diagrammed with ovals, and the operations that move the system between states are diagrammed as arrows.

The Names of the States The names of the states of the machine are represented in Z using the free type *StateCounter*. In this case, *StateCounter* is defined to be a type that is a set of five distinct variables named: *SMInitialState*, *SMS1State*, *SMS2State*, *SMS3State*, and *SMFinishedState*.

$$\textit{StateCounter} ::= \textit{SMInitialState} \mid \textit{SMS1State} \mid \textit{SMS2State} \mid \\ \textit{SMS3State} \mid \textit{SMFinishedState}$$

Free types are Z's method for constructing disjoint unions ([Spi89] 81). They are similar to union types in C ([KR88] 147), but are mathematically motivated ([BW90] 258). In this case, the free type definition introduces only new variables so it acts more like a C enumeration ([KR88] 39).

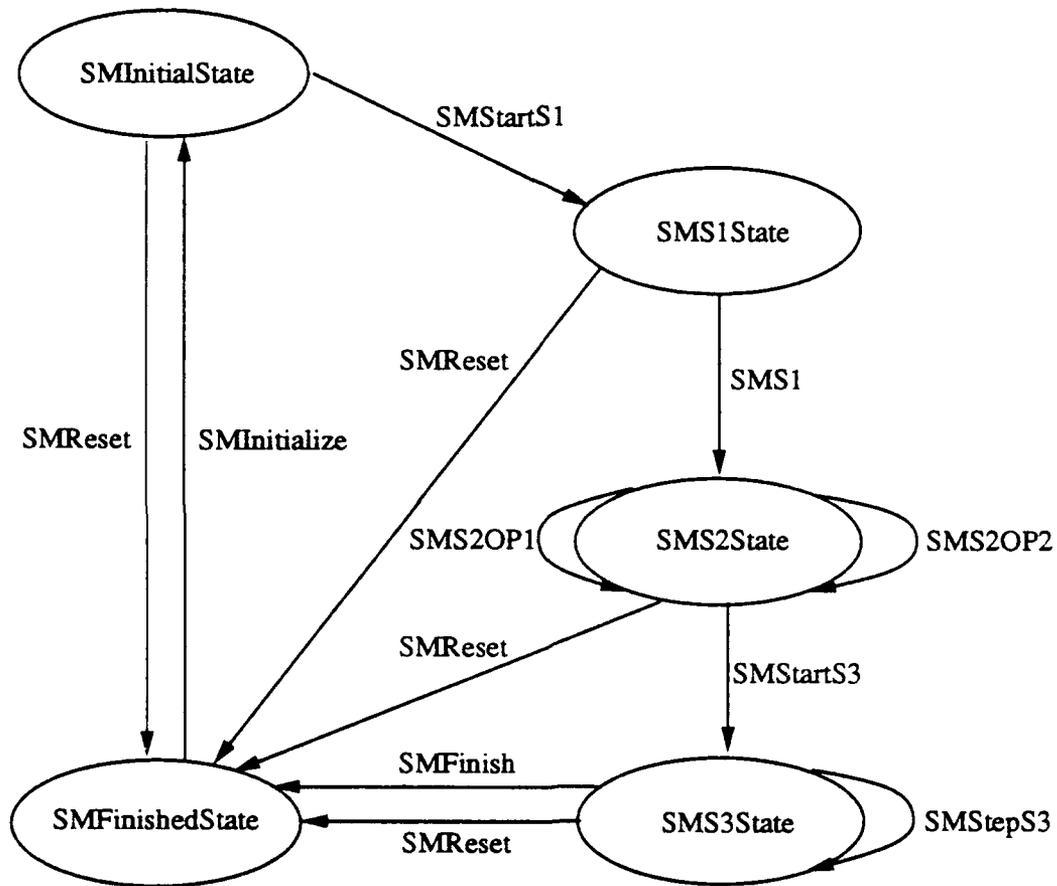


Figure 2.1: State Machine Prototype

Z defines the free type notation by expansion into an equivalent Z basic type definition ([Spi89] 50) with a set of new variables of disjoint values. In the first part of this expansion, *StateCounter* is introduced as a basic type; e.g., a set of values with no predetermined structure.

[StateCounter]

An axiomatic description is then used to introduce the five new variables. The part of the axiomatic definition above the bar is called the declaration or signature. The part below the bar is called the predicate. The declaration introduces the five new variables, typing them as elements of the set *StateCounter*. The scope of variables defined in an axiomatic description is everything after the bar in the rest of the specification. The predicate part defines predicates on the variables introduced by the declaration. In this case, the built-in Z predicate *disjoint* is applied to a sequence of the singleton sets of the variables to constrain the values to which they are bound to be distinct.

SMInitialState, SMS1State, SMS2State, SMS3State,
SMFinishedState : StateCounter

disjoint({ *SMInitialState* }, { *SMS1State* }, { *SMS2State* }, { *SMS3State* },
 { *SMFinishedState* })

The Functions of the States The five different states each perform unique functions in the prototype state machine.

1. *SMFinishedState* — when the machine is not being used, it rests in this state. The driver for this machine calls an initialize action to move it into the *SMInitialState*.
2. *SMInitialState* — this is the state from which the machine begins computation.
3. *SMS1State* — in this state the *SMS1* operation is applied to the state.
4. *SMS2State* — in this state the *SMS2OP1* or *SMS2OP2* operations are applied to the state, or *SMStartS3* is applied, moving the machine to state *S3*.
5. *SMS3State* — in this state the *SMStepS3* operation is applied to step a sub-state machine. When the sub-state machine has finished, the *SMFinish* operation ends the execution of the state machine.

Z's sequential system model uses schemas to specify almost all of the parts of a computational model. A schema can specify a data structure, like a record structure in a programming language. A schema can represent the environment

(or state) of a computational model, like the declaration of global variables in a programming language. A schema can specify predicates on an environment, like a boolean function in a program or a program invariant. A schema can also specify an operation that changes the environment, like a procedure in a programming language. Perhaps Z's greatest strength is its schema calculus. The calculus provides a rich and uniform way to combine schemas that are playing any of these roles.

The Total State of the Machine The prototype state machine's computational state is represented here using a schema named *StateSchema* ([Spi89] 51). Like an axiomatic declaration, the part above the bar is a declaration and below the bar is a predicate. The scope of variables defined in a schema is only the text from the bar to the end of the schema box. The scope of the name of the schema is all the specification past the bottom bar of the schema definition. The prototype's state schema defines only *state_counter* to remember the state the machine currently occupies. In the specification of an actual Soar state machine, the declaration's ellipses (...) would declare other variables to enrich the state model, and the predicate's ellipses would define state invariants.

```

StateSchema _____
  state_counter : StateCounter
  ...
  _____
  ...

```

Initial States for the Machine The Z convention to define a predicate on the initial states of a sequential system model is to prefix the name of the state schema with "*Init*" ([Spi89] 131). Thus the *InitStateSchema* defines a predicate which initial states of the prototype state machine must satisfy. The declaration part of this schema references the schema *StateSchema*. Z defines this reference to mean that the new schema contains all of the declarations of the referenced schema and all of the referenced schemas's predicates are conjoined with the new schema's predicate ([Spi89] 53).

```

InitStateSchema _____
  StateSchema
  _____
  state_counter = SMFinishedState
  ...
  _____

```

The expanded definition of *InitStateSchema*, shown below, is that of a schema with the same signature as the state schema, but the state counter is constrained to start in the finished state. Again, the ellipses is used to signify that the initial state schemas of state machines might contain other restrictions.

```

InitStateSchema
state_counter : StateCounter
...
state_counter = SMFinishedState
...

```

The General Form of Operations After defining the form for the state of our sequential system model and constrained its initial values, we must define operations that move the system through its state space. Operations that change the state of the sequential system are called sequential operations ([Spi89] 130). A sequential operation is a relation on pairs of states. Z uses relations instead of functions to allow the specification of non-deterministic operations. The specification takes good advantage of this: for example, decide non-deterministically selects between indifferent operator candidates by using a non-deterministic sequential operation schema.

For a schema to define a relation between two states it must have a way to reference the component variables of both states. Z's technique is to construct a schema that references both the state schema and a copy of the state schema in which all of the components have been decorated (post appended) with a single quote (') ([Spi89] 32). The undecorated state variables refer to the start state and the decorated components refer to the resultant state. For example, the *SMOperation* schema defines all of the components of *StateSchema* and a decorated copy of the components.

```

SMOperation
state_counter, state_counter' : StateCounter

```

Decorating all of the components of a schema is common enough that decorating the schema name is defined to decorate all of its components. For example, *SMOperation* could be defined by referencing both a decorated and an undecorated copy of the state schema.

```

SMOperation _____
  StateSchema
  StateSchema'
  ...

```

However, sequential operations are so common that a second short hand, Δ , is defined to mean the schema and a decorated copy of the schema.

```

SMOperation _____
   $\Delta$  StateSchema
  ...

```

The Form of a State Transition Now that we have defined the state of the system, its initial states and the general form of a state changing operation, we need to define the form of typical state machine operations. When a state machine's caller wants to initialize the system, it calls an initialization operation. Operations of this form pairs state schemas with counters in the finished state to state schemas with counters in the initial state.

```

SMInitialize _____
   $\Delta$  StateSchema
  ...
  state_counter = SMFinishedState
  state_counter' = SMInitialState
  ...

```

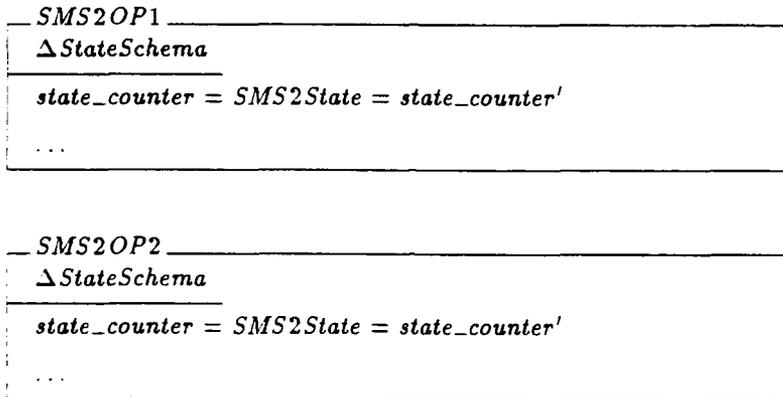
Start Transitions When a state machine can loop in a state, the operation that starts the loop by moving to the state is named with "Start".

<i>SMStartS1</i>
Δ <i>StateSchema</i>
<i>state_counter</i> = <i>SMInitialState</i>
...
<i>state_counter'</i> = <i>SMS1State</i>

Simple Transitions The *SMS1* operation provides the form for the simplest type of operation. It performs its one and only change to the state, represented by ellipses. At the same step it moves the state machine into the next state, *SMS2*.

<i>SMS1</i>
Δ <i>StateSchema</i>
<i>state_counter</i> = <i>SMS1State</i>
...
<i>state_counter'</i> = <i>SMS2State</i>

Looping and Compound Transitions SMS2 style operations are more complex than the SMS1 style operation. At each step the machine can change the sequential system state by either the SMS2OP1 operation or the SMS2OP2 operation. After the application of these operations, the machine remains in SMS2State allowing SMS2 to loop until SMS2OP1 and SMS2OP2 are both no longer applicable.



We would like to be able to define an operation that allows the state machine to take either an SMS2OP1 or an SMS2OP2 transition. The Z schema calculus ([Spi89] 74) allows the convenient combination of schemas by extending the standard logical operations of and (\wedge), or (\vee) and not (\neg) to apply to schemas. Two schemas may be combined with \wedge or \vee if their signatures give all shared variables the same types. The signature of the resulting schema contains all of the signature elements of both of the schemas and conjoins or disjuncts the predicates of the two schemas. Any schema may be negated: the resulting schema has the same signature and the negation of the predicate.

$$SMS2 \hat{=} SMS2OP1 \vee SMS2OP2$$

The SMS2 operation is defined as an "or" of the SMS2OP1 transition and the SMS2OP2 transitions. The $\hat{=}$ operation names a new schema to be the value of a schema calculus expression. This new transition satisfies the conditions of either SMS2OP1 or SMS2OP2. The SMS2 operation schema does not actually need to be defined; the machine works identically with just the SMS2OP1 and the SMS2OP2 transitions installed. (We define such transitions *only* to give the users of Soar an easy way to observe related operations: such as all the operations that leave a state, see Appendix B.) The state diagrams will not have an item representing the compound transitions because they are difficult to draw: some cases only require one arrow with two different names, but some would require a graphic merging of arrows.

In a sequential system model not all operations are applicable to all states.

The schema calculus provides an operation, called pre-condition (pre) to calculate if an operation is applicable ([Spi89] 72). The pre-condition operation maps a sequential operation to a predicate that checks if all of the conditions on the initial state of the operation are satisfied. Essentially, it tests a schema for membership in the domain of the transition relation.

SMStartS3 checks that SMS2 is no longer applicable by using the pre-condition of SMS2 . When SMS2 is no longer applicable, SMStartS3 moves the state into SMS3State . This exit transition makes SMS2State a while-do looping state: it applies SMS2 until it is no longer applicable and then exits.

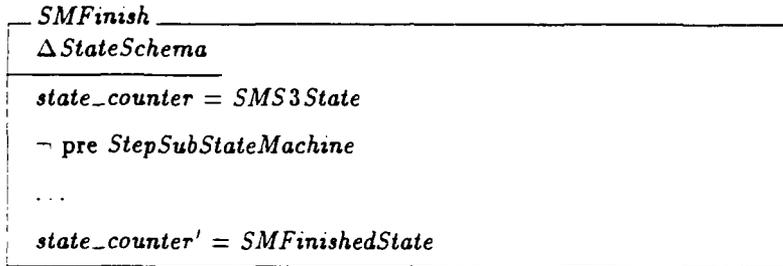
SMStartS3 $\Delta \text{StateSchema}$ $\text{state_counter} = \text{SMS2State}$ $\neg \text{pre SMS2}$ $\text{InitializeSubStateMachine}$ $\text{state_counter}' = \text{SMS3State}$
--

SubState Machine Transitions Soar is too complicated to specify with just one state machine. We need to compose machines such that one state machine can step another state machine. SMS3 prototypes a state that is implemented by a sub-state machine, named SubStateMachine . When SMStartS3 applies, it initializes the sub-state machine using $\text{InitializeSubStateMachine}$. SMStepS3 loops in the SMS3State and applies the $\text{StepSubStateMachine}$ to step the sub-state machine.

SMStepS3 $\Delta \text{StateSchema}$ $\text{state_counter} = \text{SMS3State} = \text{state_counter}'$ $\text{StepSubStateMachine}$

$$\text{SMS3} \hat{=} \text{SMStartS3} \vee \text{SMStepS3}$$

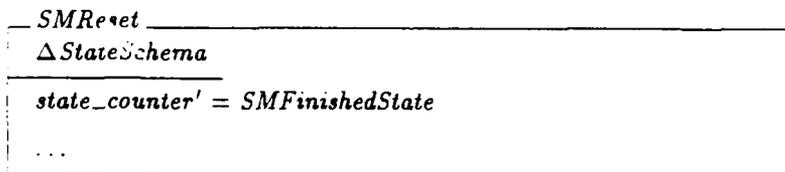
SMFinish checks that the sub-state machine is finished by checking the precondition of *StepSubStateMachine*, and moves the state machine into its *SMFinishedState*.



Stepping a Machine The caller of a state machine requires a single compound operation to step a sub-machine. *SMStep* is a prototype of an operation to step an entire machine forward one step. The stepper allows the machine to move through one step of any operation except its initialization.

$$SMStep \hat{=} SMStartS1 \vee SMS1 \vee SMS2 \vee SMS3 \vee SMFinish$$

The *SMReset* operation resets the state machine to the finished state. It is designed to be called when the state machine's execution is interrupted, and would be drawn as an arrow from every state to the finished state. As it is also cumbersome to draw this, it will not be shown in the other state machine diagrams. In an implementation, *SMReset* is designed to take actions to recover the state of the implementation that would otherwise be lost by an initialize operation.



The prototype state machine provides a convention for representing complicated control, a single step at a time, using Z's weak state relation sequential operation semantics. The convention provides templates for:

- initialization
- start transitions
- simple transitions
- looping and compound transitions

- substate machine transitions
- finishing transitions and
- machine stepping transitions.

The specification defines each piece of Soar's control flow in a state machine, and composes them into a single operation that moves Soar one atomic control step.

Chapter 3

Base Symbol Structures

This chapter defines the base sets of symbols that Soar uses to construct its knowledge representations. Soar uses two base set of symbols: Σ and the *SpecialSymbol* set. Σ is a set of base general symbols for knowledge representation. Soar composes these symbols to construct all of its knowledge structures. *SpecialSymbols* are the reserved words of Soar. They are symbols with meanings specific to the operation of Soar.

$\{\Sigma, \textit{SpecialSymbol}\}$

Soar partitions Σ into the sets of variable symbols, identifier symbols, and constant symbols. The three subsets of Σ are typed as elements of $\mathcal{P}\Sigma$, the set of all subsets of Σ ([Spi89] 27). The Z partition predicate relates a sequence of sets to a set if and only if the sets in the sequence are disjoint and their union is equal to the set ([Spi89] 125).

$\frac{\textit{Variable}, \textit{Identifier}, \textit{Constant} : \mathcal{P}\Sigma}{(\textit{Variable}, \textit{Identifier}, \textit{Constant}) \text{ partition } \Sigma}$

As many definitions allow either variables or constants in their domain we define *VariableOrConstant* for brevity.

$\frac{\textit{VariableOrConstant} : \mathcal{P}\Sigma}{\textit{VariableOrConstant} = \textit{Variable} \cup \textit{Constant}}$

Also many definitions use either an identifier or a constant in their domain. Soar users commonly refer to these as *Symbols*.

$\frac{\textit{Symbol} : \mathcal{P}\Sigma}{\textit{Symbol} = \textit{Identifier} \cup \textit{Constant}}$

The *SpecialSymbols* are partitioned into *PreferenceSymbols* and *RelationSymbols*. The *PreferenceSymbols* appear in the preference slot of preferences. The *RelationSymbols* appear in the syntax of relational tests inside of productions. Typographically, we use single quoted variable names to emphasize that these are constants of Soar, but to Z they are just Z variables. As the symbol '=' is Soar's name for both the equality relation and the indifferent preference, to we distinguish them with subscripts of "r" for relation and "p" for preference.

```

PreferenceSymbol, RelationSymbol : P SpecialSymbol
'@', '!', '+', '-', ':', '>', '<', '=p', '&' : SpecialSymbol
'=R', '<>', '<=>' : SpecialSymbol

```

(*PreferenceSymbol, RelationSymbol*) partition *SpecialSymbol*

```

RelationSymbol = { '=R', '<>', '<=>' }
PreferenceSymbol =
{ '@', '!', '+', '-', ':', '>', '<', '=p', '&' }

```

The decision procedure's input language is the set of preferences. Preferences describe properties of the structure of Soar's short term memory. The preferences may describe a property of a single object in memory, called unary, or may describe a property of two objects in memory, called binary. The sets *UnaryPreferenceSymbol* and *BinaryPreferenceSymbol* partition the preference symbols into those that can appear in unary preferences and those that can appear in binary preferences. All of the binary preference symbols are also unary, but some of the unary preferences (such as '-') can not appear in binary preferences.

```

UnaryPreferenceSymbol, BinaryPreferenceSymbol : P PreferenceSymbol

```

```

UnaryPreferenceSymbol =
{ '@', '!', '+', '-', ':', '>', '<', '=p', '&' }
BinaryPreferenceSymbol = { '>', '<', '&', '=p' }

```

The specification of the Soar architecture references 17 constant symbols. The *disjoint* predicate ensures that the variables all take no distinct values from the set of constants. No specific identifiers or variables are defined because Soar uses these as placeholders that describe the structure of its knowledge representations, but not the specific content.

```
'NIL', 'T', 'GOAL', 'PROBLEM-SPACE', 'STATE', 'OPERATOR',  
'IMPASSE', 'OBJECT', 'ITEM', 'ATTRIBUTE', 'CHOICES',  
'CONSTRAINT-FAILURE', 'CONFLICT', 'TIE', 'NO-CHANGE',  
'NONE', 'MULTIPLE', 'QUIESCENCE', 'TYPE' : Constant
```

```
disjoint ({ 'NIL' }, { 'T' }, { 'GOAL' }, { 'PROBLEM-SPACE' },  
          { 'STATE' }, { 'OPERATOR' }, { 'IMPASSE' },  
          { 'OBJECT' }, { 'ITEM' }, { 'ATTRIBUTE' },  
          { 'CHOICES' }, { 'CONSTRAINT-FAILURE' }, { 'CONFLICT' },  
          { 'TIE' }, { 'NO-CHANGE' }, { 'NONE' },  
          { 'MULTIPLE' }, { 'QUIESCENCE' }, { 'T' },  
          { 'TYPE' })
```

Chapter 4

Temporary Memory and Production Memory

This chapter specifies the elements of two of Soar's memories: temporary memory (TM) and production memory (PM). Temporary memory holds TMEs, a disjoint union type constructed from *Preferences* and *WMEs*. The preferences are also stored in preference memory (PM), and the working memory elements in working memory (WM). Temporary memory is the union of preference and working memory. Production memory holds productions that match against working memory, instantiate, fire and retract to modify preference memory.

4.1 Preferences

Soar's preferences are the input language of the decision procedure. Decide reads preference memory to determine how to change working memory. Soar has two types of preferences: unary and binary.

A Unary preference describe a property of an identifier, attribute, value triple. The most common property is that a value is acceptable for installation in WM for an identifier, attribute combination, called a slot.

```
— UnaryPreference —————  
id : Identifier  
attribute, value : Symbol  
preference : UnaryPreferenceSymbol
```

We must de-structure the data elements contained in top level memories to specify general data invariants of Soar, such as the set of all identifiers currently in use anywhere within Soar. *Components* mappings are used to de-structure an element of some type into the set of its basic symbols.

The *UnaryPreferencesComponents* mapping extracts the identifier, attribute and value symbols from a unary preference.

$$\begin{array}{l} \text{UnaryPreferencesComponents} : \text{UnaryPreference} \rightarrow \mathcal{P} \Sigma \\ \hline \forall u : \text{UnaryPreference} \bullet \\ \text{UnaryPreferencesComponents}(u) = \{u.\text{id}, u.\text{attribute}, u.\text{value}\} \end{array}$$

Binary preferences describe a property of two values for a slot to Decide. A common property described is that the identifier, attribute and value triple is a better choice for installation in working memory than the identifier, attribute and referent triple.

$$\begin{array}{l} \text{BinaryPreference} \\ \hline \text{id} : \text{Identifier} \\ \text{attribute, value} : \text{Symbol} \\ \text{preference} : \text{BinaryPreferenceSymbol} \\ \text{referent} : \text{Symbol} \end{array}$$

The *BinaryPreferencesComponents* mapping extracts the identifier, attribute, value and referent symbols from a preference.

$$\begin{array}{l} \text{BinaryPreferencesComponents} : \text{BinaryPreference} \rightarrow \mathcal{P} \Sigma \\ \hline \forall b : \text{BinaryPreference} \bullet \\ \text{BinaryPreferencesComponents}(b) = \\ \{b.\text{id}, b.\text{attribute}, b.\text{value}, b.\text{referent}\} \end{array}$$

The preference free type constructs the disjoint union of unary preferences and binary preferences. The *UnaryP* and *BinaryP* total injections ([Spi89: 106]) are used to construct a *Preference* from a unary or binary preference.

$$\text{Preference} ::= \text{UnaryP} \langle \langle \text{UnaryPreference} \rangle \rangle \mid \text{BinaryP} \langle \langle \text{BinaryPreference} \rangle \rangle$$

Given a preference, you can check if it is constructed from a unary preference or a binary preference by testing it for membership in the range of the injectors $p \in \text{ran } \textit{UnaryP}$ or $p \in \text{ran } \textit{BinaryP}$ ([Spi89] 9). The inverse functions, \textit{UnaryP}^\sim or $\textit{BinaryP}^\sim$ can then be used to unpack a preference into its unary or binary preference ([Spi89] 101). *PreferencesComponents* de-structures a preference using this technique.

$$\begin{array}{l} \textit{PreferencesComponents} : \textit{Preference} \rightarrow \mathbb{P} \textit{Symbol} \\ \forall p : \textit{Preference} \bullet \\ ((p \in \text{ran } \textit{UnaryP} \Rightarrow \\ \quad \textit{PreferencesComponents}(p) = \\ \quad \quad \textit{UnaryPreferencesComponents}(\textit{UnaryP}^\sim(p)) \wedge \\ (p \in \text{ran } \textit{BinaryP} \Rightarrow \\ \quad \textit{PreferencesComponents}(p) = \\ \quad \quad \textit{BinaryPreferencesComponents}(\textit{BinaryP}^\sim(p)))) \end{array}$$

The specification frequently creates preferences, so two functions are provided to map the components of preferences to a preference. The μ operator acts much like a programming language's let construct and the θ operator instantiates a schema using the values for its components provided in the current scope ([Spi89] 61, 64).

$$\begin{array}{l} \textit{makeUnaryPreference} : \\ \quad \textit{Identifier} \times \textit{Symbol} \times \textit{Symbol} \times \textit{UnaryPreferenceSymbol} \\ \quad \rightarrow \textit{Preference} \\ \textit{makeBinaryPreference} : \\ \quad \textit{Identifier} \times \textit{Symbol} \times \textit{Symbol} \times \textit{BinaryPreferenceSymbol} \times \textit{Symbol} \\ \quad \rightarrow \textit{Preference} \\ \forall i : \textit{Identifier}; a, v : \textit{Symbol}; p : \textit{UnaryPreferenceSymbol} \bullet \\ \quad \textit{makeUnaryPreference}(i, a, v, p) = \\ \quad \quad \textit{UnaryP}((\mu \textit{UnaryPreference} | \\ \quad \quad \quad id = i \wedge attribute = a \wedge value = v \wedge preference = p \bullet \\ \quad \quad \quad \theta \textit{UnaryPreference})) \\ \forall i : \textit{Identifier}; a, v, r : \textit{Symbol}; p : \textit{BinaryPreferenceSymbol} \bullet \\ \quad \textit{makeBinaryPreference}(i, a, v, p, r) = \\ \quad \quad \textit{BinaryP}((\mu \textit{BinaryPreference} : \\ \quad \quad \quad id = i \wedge attribute = a \wedge value = v \wedge \\ \quad \quad \quad preference = p \wedge referent = r \bullet \\ \quad \quad \quad \theta \textit{BinaryPreference})) \end{array}$$

PreferencesId is defined to allow easy access to the identifier slot of both unary and binary preferences.

$$\text{PreferencesId} : \text{Preference} \rightarrow \text{Identifier}$$

$$\forall p : \text{Preference} \bullet \\ ((p \in \text{ran UnaryP} \Rightarrow \text{PreferencesId}(p) = (\text{UnaryP}^{\sim}(p)).id) \wedge \\ (p \in \text{ran BinaryP} \Rightarrow \text{PreferencesId}(p) = (\text{BinaryP}^{\sim}(p)).id))$$

The architecture specifies that only preferences which are accessible from the most recent goal, or any impasse object, should remain in preference memory. We start the formal definition of accessible with the concept of connecting identifiers over a universe set of preferences. *Connected_{Preference}* is a mapping that takes a set of preferences to a relation between identifiers. Two identifiers are related over a set of preferences if there exists a preference that holds the first identifier in its id slot and the second in its value or referent slot.

$$\text{Connected}_{\text{Preference}} : \mathcal{P} \text{Preference} \rightarrow (\text{Identifier} \leftrightarrow \text{Identifier})$$

$$\forall P : \mathcal{P} \text{Preference}; l, r : \text{Identifier} \bullet \\ (l, r) \in (\text{Connected}_{\text{Preference}}(P)) \Leftrightarrow \\ (\exists p : P \bullet \\ ((p \in \text{ran UnaryP} \Rightarrow \\ (\text{UnaryP}^{\sim}(p)).id = l \wedge (\text{UnaryP}^{\sim}(p)).value = r) \wedge \\ (p \in \text{ran BinaryP} \Rightarrow \\ (\exists b : \text{BinaryPreference} \mid b = \text{BinaryP}^{\sim}(p) \bullet \\ b.id = l \wedge (b.value = r \vee b.referent = r))))))$$

The connection formed is roughly analogous to a directed edge in a graph; here each unary preference would be viewed as an attribute labeled edge from id to value, and each binary preference would be viewed as constructing two attribute labeled edges from id to value and from id to referent. However, the correspondence is not very natural because the node set of a graph of preferences would also have to contain all of the constants. Soar allows only identifiers to be augmented with a preference (e.g., appear in the id slot of a preference) and so the graph would require the additional constraint that only identifier nodes have out edges.

$TCI_{Preference}$ extends the concept of a single preference connecting pairs of identifiers to a path of preferences connecting sets of identifiers. The $*$ operation calculates the reflexive transitive closure of the relation $Connected_{Preference}^{(P)}$ ([Spi89] 102). The closure is then imaged ([Spi89] 101) over all of the identifiers in the start set. This produces an operation that finds all of the identifiers connected to any identifier in the start set through a path of preferences. The syntax $_(_)$ is used to denote the image of a function on a set of arguments ([Spi89] 101). The image of a function on a set, S , is the set of the function applied to all of the elements of S .

$$\underline{TCI_{Preference}} : \mathbb{P} Preference \rightarrow (\mathbb{P} Identifier \rightarrow \mathbb{P} Identifier)$$

$$\forall P : \mathbb{P} Preference; I : \mathbb{P} Identifier \bullet$$

$$(TCI_{Preference}(P))(I) = ((Connected_{Preference}(P))^*)(I)$$

$OffId_{Preference}$ extracts from a set of preferences those preferences that hold in their id component any element of the given set of identifiers.

$$\underline{OffId_{Preference}} : \mathbb{P} Identifier \times \mathbb{P} Preference \rightarrow \mathbb{P} Preference$$

$$\forall I : \mathbb{P} Identifier; P : \mathbb{P} Preference \bullet$$

$$OffId_{Preference}(I, P) = \{p : P \mid PreferencesId(p) \in I\}$$

4.2 Working Memory Elements

Soar's working memory is a set of elements named working memory elements. Working memory elements have an id component, and an attribute and value components. They are of two types: those that are a copy of an acceptable preference for a context slot and those that are not. The *context_acceptable_preference* component holds a flag of the free type, *YesOrNo* to distinguish these two types.

$$YesOrNo ::= Yes \mid No$$

WME

id : Identifier

attribute, value : Symbol

context_acceptable_preference : YesOrNo

The components of a working memory element is the set of its identifier, attribute and value.

$$\begin{array}{l} \hline \text{WMEsComponents} : WME \rightarrow \mathbb{P} \Sigma \\ \hline \forall w : WME \bullet \text{WMEsComponents}(w) = \{w.id, w.attribute, w.value\} \end{array}$$

The *makeWME* constructor is defined to simplify the instantiation of working memory element schemas.

$$\begin{array}{l} \hline \text{makeWME} : Identifier \times Symbol \times Symbol \times YesOrNo \rightarrow WME \\ \hline \forall i : Identifier; a, v : Symbol; c : YesOrNo \bullet \\ \text{makeWME}(i, a, v, c) = \\ (\mu WME \mid id = i \wedge attribute = a \wedge value = v \wedge \\ \text{context_acceptable_preference} = c \bullet \\ \theta WME) \end{array}$$

ConnectedWME defines the notion of connectedness of identifiers across a single working memory element from a universe set of working memory elements.

$$\begin{array}{l} \hline \text{ConnectedWME} : \mathbb{P} WME \rightarrow (Identifier \leftrightarrow Identifier) \\ \hline \forall W : \mathbb{P} WME; l, r : Identifier \bullet \\ (l, r) \in (\text{ConnectedWME}(W)) \Leftrightarrow \\ (\exists w : W \bullet w.id = l \wedge w.value = r) \end{array}$$

TCI_{WME} extends the connectedness definition, in the same way *TCI_{Preference}* did, to include a set of identifiers and any path of wmes.

$$\begin{array}{l} \hline \text{TCI}_{WME} : \mathbb{P} WME \rightarrow (\mathbb{P} Identifier \rightarrow \mathbb{P} Identifier) \\ \hline \forall W : \mathbb{P} WME; I : \mathbb{P} Identifier \bullet \\ (\text{TCI}_{WME}(W))(I) = ((\text{ConnectedWME}(W))^*)(I) \end{array}$$

OfIdWME extracts from a set of working memory elements, the ones that contain any element of a given set of identifiers in their id component.

$$\begin{array}{l} \hline \text{OfIdWME} : \mathbb{P} Identifier \times \mathbb{P} WME \rightarrow \mathbb{P} WME \\ \hline \forall I : \mathbb{P} Identifier; W : \mathbb{P} WME \bullet \\ \text{OfIdWME}(I, W) = \{w : W \mid w.id \in I\} \end{array}$$

4.3 Temporary Memory Elements

Temporary memory is almost a disjoint union of working memory elements and preferences. Elements of temporary memory are of type TME , a disjoint union type of $Preference$ and WME . Temporary memory is not exactly a disjoint union of working memory and preference memory. The one distinction is that the context acceptable preferences from preference memory are copied into working memory so that they can be matched by productions (see Chapter 5).

$$TME ::= PreferenceTME \langle\langle Preference \rangle\rangle \mid WMETME \langle\langle WME \rangle\rangle$$

A TME 's components are defined to be the components of the underlying preference or working memory element.

$$\begin{array}{|l} \hline TMEsComponents : TME \rightarrow \mathcal{P} \Sigma \\ \hline \forall t : TME \bullet \\ ((t \in \text{ran } PreferenceTME \Rightarrow \\ TMEsComponents(t) = \\ PreferencesComponents(PreferenceTME \sim (t))) \wedge \\ (t \in \text{ran } WMETME \Rightarrow \\ TMEsComponents(t) = \\ WMEsComponents(WMETME \sim (t)))) \\ \hline \end{array}$$

$Connected_{TME}$ defines connectedness of identifiers through a single temporary memory element over a universe of temporary memory elements.

$$\begin{array}{|l} \hline Connected_{TME} : \mathcal{P} TME \rightarrow (Identifier \leftrightarrow Identifier) \\ \hline \forall T : \mathcal{P} TME; l, r : Identifier \bullet \\ (l, r) \in Connected_{TME}(T) \Leftrightarrow \\ (l, r) \in \\ Connected_{WME}(WMETME \sim (T \cap \text{ran } WMETME)) \cup \\ Connected_{Preference}(PreferenceTME \sim (T \cap \text{ran } PreferenceTME)) \\ \hline \end{array}$$

TCI_{TME} extends the notion of connectedness to sets of identifiers over paths of temporary memory elements.

$$\begin{array}{|l} \hline TCI_{TME} : \mathcal{P} TME \rightarrow (\mathcal{P} Identifier \rightarrow \mathcal{P} Identifier) \\ \hline \forall T : \mathcal{P} TME; I : \mathcal{P} Identifier \bullet \\ (TCI_{TME}(T))(I) = ((Connected_{TME}(T))^*(I)) \\ \hline \end{array}$$

$OfId_{TME}$ extracts the subset of temporary memory elements that hold one of the given identifiers in their id slot.

$$\underline{OfId_{TME} : \mathcal{P} Identifier \times \mathcal{P} TME \rightarrow \mathcal{P} TME}$$

$$\forall I : \mathcal{P} Identifier; T : \mathcal{P} TME \bullet$$

$$OfId_{TME}(I, T) = \\ WMETME(OfId_{WME}(I, WMETME \sim (T \cap \text{ran } WMETME))) \cup \\ PreferenceTME(OfId_{Preference}(I, PreferenceTME \sim (T \cap \text{ran } PreferenceTME)))$$

4.4 Makes

The process of recall from long term memory is modeled by the matching of productions. When productions match, their right hand sides are instantiated to find preferences to add to preference memory. The right hand side is a set of makes. The makes form a language that describes how productions create new preferences, given the variable bindings from the left hand side of the production match.

The data type schema, $UnaryMake$ is instantiated to create a unary preference. Makes must contain a variable in the identifier slot, but can contain either a constant or a variable in the attribute and value fields.

$$\underline{UnaryMake}$$

$$id : Variable$$

$$attribute, value : VariableOrConstant$$

$$preference : UnaryPreferenceSymbol$$

A unary make's components is the set of its identifier, attribute and value.

$$\underline{UnaryMakesComponents} : UnaryMake \rightarrow \mathcal{P} \Sigma$$

$$\forall u : UnaryMake \bullet$$

$$UnaryMakesComponents(u) = \{u.id, u.attribute, u.value\}$$

A $BinaryMake$ is instantiated to make a binary preference.

$$\underline{BinaryMake}$$

$$id : Variable$$

$$attribute, value : VariableOrConstant$$

$$preference : BinaryPreferenceSymbol$$

$$referent : VariableOrConstant$$

A binary make's components are the set of its identifier, attribute, value and referent.

$$\begin{array}{l} \hline \textit{BinaryMakesComponents} : \textit{BinaryMake} \rightarrow \mathbb{P} \Sigma \\ \forall b : \textit{BinaryMake} \bullet \\ \textit{BinaryMakesComponents}(b) = \{b.id, b.attribute, b.value, b.referent\} \end{array}$$

The free type *Make* is defined to join the phrases of the language of makes into a single phrase.

$$\textit{Make} ::= \textit{UnaryM} \langle \langle \textit{UnaryMake} \rangle \rangle \mid \textit{BinaryM} \langle \langle \textit{BinaryMake} \rangle \rangle$$

A make's components are the components of the underlying unary or binary make.

$$\begin{array}{l} \hline \textit{MakesComponents} : \textit{Make} \rightarrow \mathbb{P} \Sigma \\ \forall m : \textit{Make} \bullet \\ ((m \in \text{ran } \textit{UnaryM} \Rightarrow \\ \textit{MakesComponents}(m) = \\ \textit{UnaryMakesComponents}(\textit{UnaryM}^{\sim}(m))) \wedge \\ (m \in \text{ran } \textit{BinaryM} \Rightarrow \\ \textit{MakesComponents}(m) = \\ \textit{BinaryMakesComponents}(\textit{BinaryM}^{\sim}(m)))) \end{array}$$

makeUnaryMake and *makeBinaryMake* instantiate copies of the unary or binary make schema, and wrap them in the make free type.

makeUnaryMake :
 Identifier × VariableOrConstant × VariableOrConstant ×
 UnaryPreferenceSymbol → Make
makeBinaryMake :
 Identifier × VariableOrConstant × VariableOrConstant ×
 BinaryPreferenceSymbol × VariableOrConstant → Make

∀ *i* : Identifier; *a, v* : VariableOrConstant; *p* : UnaryPreferenceSymbol •
makeUnaryMake(*i, a, v, p*) =
 UnaryM((μ UnaryPreference |
id = *i* ∧ *attribute* = *a* ∧ *value* = *v* ∧ *preference* = *p* •
 θ *UnaryMake*))

∀ *i* : Identifier; *a, v, r* : VariableOrConstant; *p* : BinaryPreferenceSymbol •
makeBinaryMake(*i, a, v, p, r*) =
 BinaryM((μ BinaryPreference |
id = *i* ∧ *attribute* = *a* ∧ *value* = *v* ∧
preference = *p* ∧ *referent* = *r* •
 θ *BinaryMake*))

ConnectedMake defines connectedness for makes. It is parallel to the connectedness for preferences, with the exception that makes hold constants or variables in their components instead of constants or identifiers, but connectedness only flows through variables.

ConnectedMake : P Make → (Variable ↔ Variable)

∀ *M* : P Make; *l, r* : Variable •
(*l, r*) ∈ (*ConnectedMake*(*M*)) ⇔
(∃ *m* : M •
((*m* ∈ ran *UnaryM* ⇒
(∃ *u* : *UnaryMake* | *u* = *UnaryM*[~](*m*) •
u.id = *l* ∧ *u.value* = *r*)) ∧
(*m* ∈ ran *BinaryM* ⇒
(∃ *b* : *BinaryMake* | *b* = *BinaryM*[~](*m*) •
b.id = *l* ∧ (*b.value* = *r* ∨ *b.referent* = *r*))))))

TCI_{Make} extends the definition of connectedness to cover sets of variables and paths of makes.

$$\begin{array}{l} \hline TCI_{Make} : \mathbb{P} Make \rightarrow (\mathbb{P} Variable \rightarrow \mathbb{P} Variable) \\ \hline \forall M : \mathbb{P} Make; V : \mathbb{P} Variable \bullet \\ (TCI_{Make}(M))(V) = ((Connected_{Make}(M))^*)(V) \end{array}$$

$OfId_{Make}$ extracts a subset of a set of makes that holds any element from a set of variables in their identifier slot.

$$\begin{array}{l} \hline OfId_{Make} : \mathbb{P} Variable \times \mathbb{P} Make \rightarrow \mathbb{P} Make \\ \hline \forall V : \mathbb{P} Variable; M : \mathbb{P} Make \bullet \\ OfId_{Make}(V, M) = \\ \{m : M \mid \\ (m \in \text{ran } UnaryM \Rightarrow (UnaryM \sim (m)).id \in V) \vee \\ (m \in \text{ran } BinaryM \Rightarrow (BinaryM \sim (m)).id \in V)\} \end{array}$$

4.5 Tests

Productions require a language to describe how to match against the contents of working memory. The terminal phrases of this language are called *Tests* and match against the symbols found in the fields of working memory elements.

Test is defined as a free type, both to construct a disjoint union and to allow tests to recursively contain tests ([Spi89] 81-83). Three of the test types are constant, and five are constructed from arguments, one of which takes a set of tests as arguments.

$$\begin{array}{l} Test ::= BlankTest \\ | YesTest \\ | NoTest \\ | EqualityTest \langle \langle \Sigma \rangle \rangle \\ | NotTest \langle \langle \Sigma \rangle \rangle \\ | SameTypeTest \langle \langle \Sigma \rangle \rangle \\ | DisjunctiveTest \langle \langle \mathbb{P} Constant \rangle \rangle \\ | ConjunctiveTest \langle \langle \mathbb{P} Test \rangle \rangle \end{array}$$

As a production matches, it constructs a binding of variables to symbols, much like an environment in a programming language. When a test matches against a symbol, if the test is constructed using a variable it will match using the binding of the variable instead of the variable itself.

There are eight types of tests.

1. *BlankTest* — matches any symbol.
2. *YesTest* — matches against a value of type *YesOrNo* that is a *Yes*.
3. *NoTest* — matches against a value of type *YesOrNo* that is a *No*.
4. *EqualityTest* — matches against a symbol if it is the same symbol used to create the equality test, or the binding of the variable used to create the test.
5. *NotTest* — matches against a symbol, only if is not the same symbol used to create the *NotTest*, or not the same symbol as the binding of the variable used to create the test.
6. *SameTypeTest* — matches against a symbol only if the symbol and the constructing argument of the *SameTypeTest* (or its binding) are both in *Identifier* or both in *Constant*.
7. *DisjunctiveTest* — matches a symbol only if it is in the set of constants used to create the disjunctive test.
8. *ConjunctiveTest* — matches a symbol only if the symbol matches all of the set of tests that were used to create the conjunction.

Although Soar could support recursive conjunctive tests they provide no extra matching power over a single layer of conjunctions. The predicate below constrains all conjunctive tests to be constructed only from equality tests, not tests, same type tests and disjunctions.

$$\begin{aligned}
 \forall ct : \text{ran } \textit{ConjunctiveTest} \bullet \\
 \textit{ConjunctiveTest} \sim (ct) \subseteq \\
 (\text{ran } \textit{EqualityTest} \cup \text{ran } \textit{NotTest} \cup \\
 \text{ran } \textit{SameTypeTest} \cup \text{ran } \textit{DisjunctiveTest})
 \end{aligned}$$

The components of tests are collected from the arguments used to construct the test.

TestsComponents : *Test* → P Σ

$\forall t : \text{Test} \bullet$
 $((t \in \{\text{BlankTest}, \text{YesTest}, \text{NoTest}\} \Rightarrow$
 $\quad \text{TestsComponents}(t) = \emptyset) \wedge$
 $(t \in \text{ran EqualityTest} \Rightarrow$
 $\quad \text{TestsComponents}(t) = \{\text{EqualityTest}^{\sim}(t)\}) \wedge$
 $(t \in \text{ran NotTest} \Rightarrow$
 $\quad \text{TestsComponents}(t) = \{\text{NotTest}^{\sim}(t)\}) \wedge$
 $(t \in \text{ran SameTypeTest} \Rightarrow$
 $\quad \text{TestsComponents}(t) = \{\text{SameTypeTest}^{\sim}(t)\}) \wedge$
 $(t \in \text{ran DisjunctiveTest} \Rightarrow$
 $\quad \text{TestsComponents}(t) = \text{DisjunctiveTest}^{\sim}(t)) \wedge$
 $(t \in \text{ran ConjunctiveTest} \Rightarrow$
 $\quad \text{TestsComponents}(t) = \bigcup(\text{TestsComponents}(\text{ConjunctiveTest}^{\sim}(t))))$

Soar's conditions have a notion of a positive versus a negative test. A positive test is one that checks for equality with a constant or a variable's binding, and a negative test is one that does not. The *TestsPositiveComponents* mapping is defined to extract the set of constants and variables that a test positively checks.

TestsPositiveComponents : *Test* → P Symbol

$\forall t : \text{Test} \bullet$
 $((t \in \{\text{BlankTest}, \text{YesTest}, \text{NoTest}\} \cup$
 $\quad \text{ran SameTypeTest} \cup \text{ran NotTest} \Rightarrow$
 $\quad \text{TestsPositiveComponents}(t) = \emptyset) \wedge$
 $(t \in \text{ran EqualityTest} \Rightarrow$
 $\quad \text{TestsPositiveComponents}(t) = \{\text{EqualityTest}^{\sim}(t)\}) \wedge$
 $(t \in \text{ran DisjunctiveTest} \Rightarrow$
 $\quad \text{TestsPositiveComponents}(t) = \text{DisjunctiveTest}^{\sim}(t)) \wedge$
 $(t \in \text{ran ConjunctiveTest} \Rightarrow$
 $\quad \text{TestsPositiveComponents}(t) =$
 $\quad \bigcup(\text{TestsPositiveComponents}(\text{ConjunctiveTest}^{\sim}(t))))$

While describing operations on tests, it is often necessary to describe one action for conjunctive tests and another action for all other types of tests. The abbreviation definition ([Spi89] 52) below defines a new type set, *SimpleTest*, that simplifies checking whether a test is not conjunctive.

$$\text{SimpleTest} == (\text{Test} \setminus \text{ran ConjunctiveTest})$$

4.6 WME Tests

Tests are composed in structures, called a *WMETest*, that tests against a working memory element. A *WMETest* matches a working memory element if and only if all of the working memory element test's tests match the corresponding fields of the working memory element.

$$\begin{array}{l} \text{--- } \underline{\text{WMETest}} \text{ ---} \\ \text{id, attribute : Test} \setminus \{\text{BlankTest, YesTest, NoTest}\} \\ \text{value : Test} \setminus \{\text{YesTest, NoTest}\} \\ \text{context_acceptable_preference : \{YesTest, NoTest\}} \end{array}$$

makeWMETest instantiates a working memory element test given its four component tests.

$$\begin{array}{l} \text{makeWMETest : Test} \setminus \{\text{BlankTest, YesTest, NoTest}\} \times \\ \text{Test} \setminus \{\text{YesTest, NoTest}\} \times \\ \text{Test} \setminus \{\text{YesTest, NoTest}\} \times \\ \{\text{YesTest, NoTest}\} \times \{\text{YesTest, NoTest}\} \rightarrow \text{WMETest} \\ \text{---} \\ \forall \text{id : Test} \setminus \{\text{BlankTest, YesTest, NoTest}\}; \\ \text{attribute, value : Test} \setminus \{\text{YesTest, NoTest}\}; \\ \text{context_acceptable_preference : \{YesTest, NoTest\}} \bullet \\ \text{makeWMETest(id, attribute, value, context_acceptable_preference) =} \\ (\mu \text{id : \{id\}; attribute : \{attribute\}; value : \{value\};} \\ \text{context_acceptable_preference : \{context_acceptable_preference\}} \bullet \\ \theta \text{WMETest}) \end{array}$$

The components of a working memory element test are the components of its id, attribute and value fields.

$$\begin{array}{l} \text{--- } \underline{\text{WMETestsComponents}} \text{ ---} \\ \text{WMETestsComponents : WMETest} \rightarrow \mathbb{P} \text{ Symbol} \\ \text{---} \\ \forall \text{wt : WMETest} \bullet \\ \text{WMETestsComponents(wt) =} \\ \bigcup (\text{TestsComponents}(\{\text{wt.id, wt.attribute, wt.value}\})) \end{array}$$

WME Tests Positive Components extends the concept of a positive component to the level of a working memory element test.

$$\begin{array}{l}
 \hline
 \text{WME Tests Positive Components} : \text{WME Test} \rightarrow \mathbb{P} \text{ Symbol} \\
 \hline
 \forall wt : \text{WME Test} \bullet \\
 \quad \text{WME Tests Positive Components}(wt) = \\
 \quad \quad \cup(\text{Tests Positive Components}(\{wt.id, wt.attribute, wt.value\}))
 \end{array}$$

Two working memory element tests are connected if either of their identifier tests share positive components, or the first test's value test shares components with the second one's id test.

$$\begin{array}{l}
 \hline
 \text{Connected}_{\text{WME Test}} - : \text{WME Test} \rightarrow \text{WME Test} \\
 \hline
 \forall wt_1, wt_2 : \text{WME Test} \bullet \\
 \quad wt_1 \text{ Connected}_{\text{WME Test}} wt_2 \Leftrightarrow \\
 \quad \quad ((\text{Tests Positive Components}(wt_1.id) \cap \\
 \quad \quad \quad \text{Tests Positive Components}(wt_2.id)) \neq \emptyset \vee \\
 \quad \quad (\text{Tests Positive Components}(wt_1.value) \cap \\
 \quad \quad \quad \text{Tests Positive Components}(wt_2.id)) \neq \emptyset)
 \end{array}$$

4.7 Conditions

Soar's productions check working memory for the existence of elements matching working memory element tests, but they must also check if working memory does not contain matching patterns of elements. Working memory element tests are bundled into structures, called *Conditions*, that match on the existence and non-existence of patterns of elements.

Positive Conditions match when there exists a working memory element that their *WME Test* matches. *Negative Condition* match the non-existence of a element that matches their *WME Test*. *Negative Conjunctive Conditions* match when all of their conditions fail to simultaneously match working memory.

$$\begin{array}{l}
 \text{Condition} ::= \text{Positive Condition} \langle \langle \text{WME Test} \rangle \rangle \\
 \quad | \text{Negative Condition} \langle \langle \text{WME Test} \rangle \rangle \\
 \quad | \text{Negative Conjunctive Condition} \langle \langle \mathbb{P}_1 \text{ Condition} \rangle \rangle
 \end{array}$$

Conjunctive negations are restricted to contain at least one positive condition.

$$\begin{aligned} &\forall c : \text{ran } \textit{NegativeConjunctiveCondition} \bullet \\ &\quad \exists cs : \mathbb{P} \textit{Condition} \mid cs = \textit{NegativeConjunctiveCondition} \sim (c) \bullet \\ &\quad \text{ran } \textit{PositiveCondition} \cap cs \neq \emptyset \end{aligned}$$

While specifying operations on conditions cased by type, it is convenient to have a sub-type of conditions, *SimpleCondition*, that excludes the negative conjunctive conditions.

$$\textit{SimpleCondition} == \text{ran } \textit{PositiveCondition} \cup \text{ran } \textit{NegativeCondition}$$

ConditionsComponents extracts the components of a condition from its working memory element tests or recursively from its embedded conditions.

$$\textit{ConditionsComponents} : \textit{Condition} \rightarrow \mathbb{P} \textit{Symbol}$$

$$\begin{aligned} &\forall c : \textit{Condition} \bullet \\ &\quad ((c \in \text{ran } \textit{PositiveCondition} \Rightarrow \\ &\quad \quad \textit{ConditionsComponents}(c) = \\ &\quad \quad \quad \textit{WMETestsComponents}(\textit{PositiveCondition} \sim (c))) \wedge \\ &\quad (c \in \text{ran } \textit{NegativeCondition} \Rightarrow \\ &\quad \quad \textit{ConditionsComponents}(c) = \\ &\quad \quad \quad \textit{WMETestsComponents}(\textit{NegativeCondition} \sim (c))) \wedge \\ &\quad (c \in \text{ran } \textit{NegativeConjunctiveCondition} \Rightarrow \\ &\quad \quad \textit{ConditionsComponents}(c) = \\ &\quad \quad \quad \bigcup (\textit{ConditionsComponents} \{ \textit{NegativeConjunctiveCondition} \sim (c) \}))) \end{aligned}$$

The positive components of a condition are defined to be all of the positive components of its working memory element test or embedded conditions.

$$\textit{ConditionsPositiveComponents} : \textit{Condition} \rightarrow \mathbb{P} \textit{Symbol}$$

$$\begin{aligned} &\forall c : \textit{Condition} \bullet \\ &\quad ((c \in \text{ran } \textit{PositiveCondition} \Rightarrow \\ &\quad \quad \textit{ConditionsPositiveComponents}(c) = \\ &\quad \quad \quad \textit{WMETestsPositiveComponents}(\textit{PositiveCondition} \sim (c))) \wedge \\ &\quad (c \in \text{ran } \textit{NegativeCondition} \Rightarrow \\ &\quad \quad \textit{ConditionsPositiveComponents}(c) = \\ &\quad \quad \quad \textit{WMETestsPositiveComponents}(\textit{NegativeCondition} \sim (c))) \wedge \\ &\quad (c \in \text{ran } \textit{NegativeConjunctiveCondition} \Rightarrow \\ &\quad \quad \textit{ConditionsPositiveComponents}(c) = \\ &\quad \quad \quad \bigcup (\textit{ConditionsPositiveComponents} \{ \textit{NegativeConjunctiveCondition} \sim (c) \}))) \end{aligned}$$

Soar constrains the conditions of the left hand side of the production to

be connected, and chunking's inner loops calculate the transitive closures of conditions. Unfortunately, negated conditions and conjunctive negations do not allow the definition of connectedness of conditions to be parallel with those of preferences, working memory elements, temporary memory elements and makes.

ConnectedCondition is defined to connect sets of conditions to sets of conditions, instead of connecting pairs of identifiers over a universe of conditions. This allows the definition to support Soar's use of multiple starting conditions to connect into a conjunctive negation.

Two sets of conditions, A and B, are connected if and only if B holds a positive condition, b, whose *WMETest* is connected to the *WMETest* of any positive condition in A, and A union the set of b is connected to B's other conditions. Similarly, a negated condition, b in B, is connected to A if there is a positive condition in A whose *WMETest* connects to b, and A connects to the remaining elements of B. Notice that only positive conditions in A may connect to a elements of B.

A set of conditions, A, is connected to a set of conditions, B, containing a conjunctive negation, b if and only if A connects to all of the sub-conditions of b and A union the set of b connects to B's remaining conditions. This connects any conjunctive negation whose sub-conditions are all connected to A, either directly or through a path in the sub-conditions.

ConnectedCondition - : $\mathcal{P} \text{ Condition} \leftrightarrow \mathcal{P} \text{ Condition}$

$\forall A : \mathcal{P} \text{ Condition} \bullet A \text{ ConnectedCondition} \emptyset$

$\forall A, B : \mathcal{P} \text{ Condition} \bullet$

$A \text{ ConnectedCondition} B \Leftrightarrow$

$(\exists b : B \bullet$

$(b \in \text{ran PositiveCondition} \Rightarrow$

$(\exists a : A \bullet$

$\text{PositiveCondition} \sim (a) \text{ ConnectedWMETest PositiveCondition} \sim (b))) \wedge$

$(b \in \text{ran NegativeCondition} \Rightarrow$

$(\exists a : A \bullet$

$\text{PositiveCondition} \sim (a) \text{ ConnectedWMETest NegativeCondition} \sim (b))) \wedge$

$(b \in \text{ran NegativeConjunctiveCondition} \Rightarrow$

$(A \text{ ConnectedCondition NegativeConjunctiveCondition} \sim (b))) \wedge$

$A \cup \{b\} \text{ ConnectedCondition} B \setminus \{b\}$

$TCI_{Condition}$ transitively closes a set of conditions S over a universe U . The definition constructs the closure using the maximal set under the set inclusion ordering. The image of a form of the reflexive transitive closure could be used, but it would be cumbersome due to $Connected_{Condition}$'s different signature.

$$\begin{array}{l}
 TCI_{Condition} : \mathcal{P} \text{ Condition} \times \mathcal{P} \text{ Condition} \rightarrow \mathcal{P} \text{ Condition} \\
 \hline
 \forall S, U : \mathcal{P} \text{ Condition} \bullet \\
 \quad \exists C : \mathcal{P} U \mid \\
 \quad (S \text{ Connected}_{Condition} C) \wedge \\
 \quad (\forall B : \mathcal{P} U \bullet S \text{ Connected}_{Condition} B \Rightarrow B \subseteq C) \bullet \\
 TCI_{Condition}(S, U) = C
 \end{array}$$

$OfId_{Condition}$ extracts the positive and negative conditions of a set that match any constant in their identifier. It does not extract the negative conjunctive conditions that hold the set of symbols in their identifier because this requires a full recursive definition of connectedness. Where this is required, $OfId_{Condition}$ is used in conjunction with $Connected_{Condition}$ so that it can connect into the conjunctive negations.

$$\begin{array}{l}
 OfId_{Condition} : \mathcal{P} \Sigma \times \mathcal{P} \text{ Condition} \rightarrow \mathcal{P} \text{ Condition} \\
 \hline
 \forall I : \mathcal{P} \Sigma; C : \mathcal{P} \text{ Condition} \bullet \\
 OfId_{Condition}(I, C) = \\
 \quad \{c : C \mid \exists i : I \bullet \\
 \quad (c \in \text{ran } \text{PositiveCondition} \Rightarrow \\
 \quad \quad i \in \text{TestsPositiveComponents}((\text{PositiveCondition} \sim (c)).id)) \wedge \\
 \quad (c \in \text{ran } \text{NegativeCondition} \Rightarrow \\
 \quad \quad i \in \text{TestsPositiveComponents}((\text{NegativeCondition} \sim (c)).id))\}
 \end{array}$$

4.8 SP

Productions are Soar's model of long term memories. They match against the contents of working memory, instantiate and fire to augment support memory with preferences.

Soar's definition of match requires that every variable that is tested against (with NotTest, or SameTypeTest) must be bound in some condition (with an EqualityTest). *TestedVariablesBound* is a predicate that checks that the variables a condition tests against are all in the set of bound variables

$$\begin{aligned}
 & \text{TestedVariablesBound} _ : \mathbb{P}(\text{Condition} \times \mathbb{P} \text{Variable}) \\
 & \forall \text{bound} : \mathbb{P} \text{Variable}; c : \text{Condition} \bullet \\
 & \quad \text{TestedVariablesBound}(c, \text{bound}) \Leftrightarrow \\
 & \quad (c \in \text{ran PositiveCondition} \Rightarrow \\
 & \quad \quad (\text{ConditionsComponents}(c) \cap \text{Variable}) \subseteq \text{bound}) \wedge \\
 & \quad (c \in \text{ran NegativeCondition} \Rightarrow \\
 & \quad \quad (\text{ConditionsComponents}(c) \cap \text{Variable}) \subseteq \\
 & \quad \quad (\text{bound} \cup (\text{ConditionsPositiveComponents}(c) \cap \text{Variable}))) \wedge \\
 & \quad (c \in \text{ran NegativeConjunctiveCondition} \Rightarrow \\
 & \quad \quad (\exists cs : \{\text{NegativeConjunctiveCondition}^{\sim}(c)\} \bullet \\
 & \quad \quad (\exists bv : \mathbb{P} \text{Variable} \{ \\
 & \quad \quad \quad bv = \text{bound} \cup \\
 & \quad \quad \quad (\bigcup (\text{ConditionsPositiveComponents}(cs) \cap \text{ran PositiveCondition}) \\
 & \quad \quad \quad \cap \text{Variable}) \bullet \\
 & \quad \quad (\forall c : cs \bullet \text{TestedVariablesBound}(c, bv))))))
 \end{aligned}$$

Soar's productions, SPs, have a name, a set of left hand side conditions, lhs, that describe how they match against working memory and a right hand side set of actions, rhs. The lhs must contain at least one positive condition, or the rhs would have no variables bound to use in make actions. Each condition may only test against variables that are bound with a equality test somewhere on the left hand side.

<i>SP</i>
<i>name</i> : <i>Symbol</i>
<i>lhs</i> : \mathbb{P}_1 <i>Condition</i>
<i>rhs</i> : \mathbb{P} <i>Make</i>
<i>roots</i> : \mathbb{P} <i>Variable</i>
$\exists c : lhs \bullet c \in \text{ran } \textit{PositiveCondition}$
$\forall c : lhs \bullet \textit{ConditionsComponents}(c) \cap \textit{Identifier} = \emptyset$
<i>roots</i> =
$\{v : \bigcup(\textit{ConditionsPositiveComponents}(\{lhs \cap \text{ran } \textit{PositiveCondition}\}))$ $\quad \neg (\exists c : lhs \cap \text{ran } \textit{PositiveCondition} \bullet$ $\quad \quad v \in \textit{TestsPositiveComponents}(\{\textit{PositiveCondition} \sim (c)\}.value))\}$
<i>lhs</i> = $\textit{TCICondition}(\textit{OfIdCondition}(\textit{roots}, lhs), lhs)$
$\exists \textit{boundvariables} : \mathbb{P}$ <i>Variable</i>
<i>boundvariables</i> =
$(\bigcup(\textit{ConditionsPositiveComponents}(\{lhs \cap \text{ran } \textit{PositiveCondition}\})$ $\quad \cap \textit{Variable}) \bullet$ $(\forall c : lhs \bullet \textit{TestedVariablesBound}(c, \textit{boundvariables}) \wedge$ $\textit{rhs} = \textit{OfIdMake}(\textit{TCIMake}(\textit{rhs})(\textit{boundvariables}), \textit{rhs}))$

Conditions are stored in two places in Soar: in productions they describe when and how productions match, and in chunking they record which patterns their productions matched against. Chunking's conditions may have matched against identifiers and so they are instantiated with identifiers to record the match. However, productions are constrained not to use identifiers in their match, otherwise they would be terribly specific to the set of identifiers in use (except for internal chunks).

A root of a production is an identifier that is not positively tested in any value of the positive conditions. Soar constrains these variables to match against goal or impasse identifiers. All of the conditions of the lhs must be in the transitive closure of the roots; this ensures that productions only match working memory elements which are in the transitive closure of the context stack or an impasse.

The right hand side is also constrained to be transitively closed over the bound variables of the left hand side. This ensures that the instantiations only make new preferences that are accessible from a goal or impasse.

makeSP is used to construct a new production data schema. It takes only three arguments, the name, the left hand side and the right hand side, as the roots are a derived component of the left hand side.

$$\frac{\text{makeSP} : \text{Symbol} \times \mathbb{P}_1 \text{Condition} \times \mathbb{P} \text{Make} \rightarrow \text{SP}}{\begin{array}{l} \forall \text{name} : \text{Symbol}; \text{lhs} : \mathbb{P}_1 \text{Condition}; \text{rhs} : \mathbb{P} \text{Make} \bullet \\ \exists_1 \text{roots} : \mathbb{P} \text{Variable} \mid \text{SP} \bullet \\ \text{makeSP}(\text{name}, \text{lhs}, \text{rhs}) = \theta \text{SP} \end{array}}$$

The components of productions are all of the components of its left hand side and the components of its right hand side.

$$\frac{\text{SpsComponents} : \text{SP} \rightarrow \mathbb{P} \Sigma}{\begin{array}{l} \forall \text{sp} : \text{SP} \bullet \\ \text{SpsComponents}(\text{sp}) = \\ \cup(\text{ConditionsComponents}(\text{sp.lhs}) \cup \text{MakesComponents}(\text{sp.rhs})) \end{array}}$$

Chapter 5

Recognition Memory

Recognition memory is where Soar holds, processes and acquires long term memories. Long term memories are stored as productions that match against short term memories in working memory. Recognition memory matches the productions against working memory, stores the matches, instantiates matches to produce new preferences, and calculates support for preferences and chunks.

The chapter is divided into nine sections:

- Section 5.1 Bindings — introduces bindings that map variables to values
- Section 5.2 Matching — defines how productions match against working memory
- Section 5.3 Instantiation — specifies how matches are instantiated
- Section 5.4 Goal Operations — defines some operations on Decide's goal stack for use by support
- Section 5.5 Chunking — defines chunking, Soar's learning mechanism
- Section 5.6 Slots — defines the concepts of a slot.
- Section 5.7 Recognition Memory's State — constructs the state of recognition memory
- Section 5.8 Support — specifies how instantiations support preferences and
- Section 5.9 Preference Phase Operations — defines the state machine for recognition memory's control loop, called preference phase.

5.1 Bindings

Soar's productions match against working memory using a binding of the production's variables to symbols. Bindings are similar to environments in programming languages. In Soar, bindings are modeled as partial functions from variables to symbols.

$$\textit{Binding} == \textit{Variable} \rightarrow \textit{Symbol}$$

When productions match, they check that there is no possible way to instantiate their negated conditions to allow them to match working memory. The negative conditions must not be able to add bindings for their local variables that allow them to match working memory elements. To formalize the match of negative and negative conjunctive conditions, we define the concepts of two bindings being consistent, a binding covering a set of variables, and one binding being a consistent extension of another.

Two bindings are consistent if they both bind every variable that they share to the same symbol. A binding covers a set of variables if it gives a binding to every variable in the set. A binding, A, is a consistent extension of another binding, B, if it is consistent with it and A possibly binds variables that B does not.

$$_ \textit{Consistent} _ , _ \textit{ConsistentExtension} _ : \textit{Binding} \leftrightarrow \textit{Binding}$$

$$_ \textit{Covers} _ : \textit{Binding} \leftrightarrow \mathbb{P} \Sigma$$

$$\forall a, b : \textit{Binding} \bullet$$

$$(a \textit{Consistent} b) \Leftrightarrow$$

$$(\forall v : \textit{dom} a \cap \textit{dom} b \bullet a(v) = b(v))$$

$$\forall b : \textit{Binding} \bullet \forall r : \mathbb{P} \textit{Variable} \bullet$$

$$b \textit{Covers} r \Leftrightarrow r \subseteq \textit{dom} b$$

$$\forall a, b : \textit{Binding} \bullet$$

$$a \textit{ConsistentExtension} b \Leftrightarrow$$

$$(a \textit{Consistent} b \wedge (\textit{dom} b \subseteq \textit{dom} a))$$

5.2 Matching

The core result of matching a production against working memory is a correspondence between the positive conditions of the production and the working memory elements they match. In this section, we define how a test matches a symbol, and how a working memory element test matches a working memory element. We use this to define the concept of a match to a production, and then define when a match is consistent with the contents of working memory.

5.2.1 Tests Matching Symbols

A test can match a symbol, under a binding, in one of six ways.

1. The test can be the blank test, which matches any symbol.
2. Equality tests only match a symbol, A , if their argument is A , or a variable bound to A . The \oplus notation is Z's way of function overriding ([Spi89] 108). $(a \oplus b)(x)$ is $b(x)$ if $x \in \text{dom } b$ or $a(x)$ if $x \in \text{dom } a$ and undefined otherwise. The $(\text{id } \Sigma)$ is the identity function on Σ ([Spi89] 97). The overriding of $(\text{id } \Sigma)$ with the binding, b , produces a function that returns the binding of the equality test's variable, or the equality test's constant.
3. Not tests match a symbol, A , only if their symbol or variable's binding is not equal to A .
4. Same type tests match a symbol, A , against their symbol or variable's binding, B , only if A and B are both identifiers or both constants.
5. A disjunctive test matches only constants that are elements of the disjunction's argument set.
6. Conjunctive tests match a symbol only if the symbol matches all of the conjunction's sub-tests.

$$\text{Matches}_{\text{Symbol}}^{\text{Test}} - : \mathbb{P}(\text{Binding} \times \text{Test} \times \text{Symbol})$$

$$\begin{aligned} & \forall b : \text{Binding}; t : \text{Test}; v : \text{Symbol} \bullet \\ & ((t = \text{BlankTest}) \vee \\ & (t \in \text{ran } \text{EqualityTest} \wedge v = ((\text{id } \Sigma) \oplus b)(\text{EqualityTest}^{\sim}(t))) \vee \\ & (t \in \text{ran } \text{NotTest} \wedge v \neq ((\text{id } \Sigma) \oplus b)(\text{NotTest}^{\sim}(t))) \vee \\ & (t \in \text{ran } \text{SameTypeTest} \wedge \\ & (\exists v2 : \{((\text{id } \text{Symbol}) \oplus b)(\text{SameTypeTest}^{\sim}(t))\} \bullet \\ & ((v \in \text{Identifier} \wedge v2 \in \text{Identifier}) \vee \\ & (v \in \text{Constant} \wedge v2 \in \text{Constant})))) \wedge \\ & (t \in \text{ran } \text{DisjunctiveTest} \wedge v \in \text{DisjunctiveTest}^{\sim}(t)) \vee \\ & (t \in \text{ran } \text{ConjunctiveTest} \wedge \\ & (\forall \text{subt} : \text{ConjunctiveTest}^{\sim}(t) \bullet \text{Matches}_{\text{Symbol}}^{\text{Test}}(b, \text{subt}, v)))) \end{aligned}$$

A working memory element test matches a working memory element if its id, attribute, value and context acceptable preference fields match the element's test's id, attribute, value and context acceptable preference test.

$$\text{Matches}_{WME}^{WME\text{Test}} _ : \mathcal{P}(\text{Binding} \times WME\text{Test} \times WME)$$

$$\forall b : \text{Binding}; wt : WME\text{Test}; w : WME \bullet$$

$$\text{Matches}_{WME}^{WME\text{Test}}(b, wt, w) \Leftrightarrow$$

$$(\text{Matches}_{\text{Symbol}}^{\text{Test}}(b, wt.id, w.id) \wedge$$

$$\text{Matches}_{\text{Symbol}}^{\text{Test}}(b, wt.attribute, w.attribute) \wedge$$

$$\text{Matches}_{\text{Symbol}}^{\text{Test}}(b, wt.value, w.value) \wedge$$

$$(wt.context_acceptable_preference = \text{YesTest} \Rightarrow$$

$$w.context_acceptable_preference = \text{Yes}) \wedge$$

$$(wt.context_acceptable_preference = \text{NoTest} \Rightarrow$$

$$w.context_acceptable_preference = \text{No}))$$

5.2.2 Matches

The full result of a production matching is recorded in a match schema. The match contains: the production that matched, the correspondence between its positive conditions and the working memory elements it matched, called a matching, the binding that allows this correspondence, and the set of matched working memory elements.

The matching is actually a correspondence between the positive conditions and working memory element, number pairs. As each working memory element could be added to working memory, later removed and the re-added, we identify each instance of a working memory element uniquely using a natural number tag.

<p><i>Match</i></p> <p><i>production</i> : <i>SP</i></p> <p><i>matching</i> : $\text{ran } \textit{PositiveCondition} \rightarrow (WME \times \mathbb{N})$</p> <p><i>binding</i> : <i>Binding</i></p> <p><i>lhs</i> : $\mathbb{P} WME$</p> <hr/> <p>$\text{dom } \textit{matching} = (\textit{production}.\textit{lhs} \cap \text{ran } \textit{PositiveCondition})$</p> <p>$\textit{lhs} = \text{first}(\text{ran } \textit{matching})$</p> <p>$\text{dom } \textit{binding} =$ $\bigcup (\textit{ConditionsPositiveComponents}(\textit{production}.\textit{lhs} \cap \text{ran } \textit{PositiveCondition}))$ $\cap \textit{Variable}$</p> <p>$\forall c : \text{dom } \textit{matching} \bullet$ $\text{Matches}_{WME}^{WME\text{Test}} (\textit{binding}.\textit{PositiveCondition} \sim (c), \text{first}(\textit{matching}(c)))$</p>

makeMatch constructs a new match given the production and the matching between the positive conditions and the working memory elements that they matched. The match's binding and left hand side are uniquely determined by the production and the matching, so they are not required arguments to construct a match.

<p><i>makeMatch</i> :</p> <p>$SP \times (\text{ran } \textit{PositiveCondition} \rightarrow (WME \times \mathbb{N})) \rightarrow \textit{Match}$</p> <hr/> <p>$\forall \textit{production} : SP;$ $\textit{matching} : (\text{ran } \textit{PositiveCondition} \rightarrow (WME \times \mathbb{N})) \bullet$ $\exists_1 \textit{binding} : \textit{Binding}; \textit{lhs} : \mathbb{P} WME \mid \textit{Match} \bullet$ $\text{makeMatch}(\textit{production}, \textit{matching}) = \theta \textit{Match}$</p>
--

5.2.3 Matching against Working Memory

The definition of match checks the correspondence between the positive conditions and the working memory elements that they matched, but it makes no statement that these elements must be in working memory. Nor does it require that the negative conditions do not have any matches in working memory. We rectify this by constraining how a condition matches against the contents of working working memory

$\text{Matches}_{WM}^{Condition}$ takes a binding, a condition and a set of working memory elements, and checks if the condition matches working memory. A positive condition matches working memory if there is a element in working memory that its WME_{Test} matches under the binding. A negative condition matches against working memory if there is no way to consistently extend the binding to cover the variables of the negated condition and then find an element in working memory whose test matches. A negative conjunctive condition matches working memory if there is no way to extend the binding and to match all of its sub-conditions against working memory.

$$\text{Matches}_{WM}^{Condition} : \mathcal{P}(\text{Binding} \times \text{Condition} \times \mathcal{P} \text{WME})$$

$$\forall b : \text{Binding}; pc : \text{ran PositiveCondition}; WM : \mathcal{P} \text{WME} \bullet$$

$$\text{Matches}_{WM}^{Condition} (b, pc, WM) \Leftrightarrow$$

$$(\exists w : \text{WME} \bullet$$

$$\text{Matches}_{WME}^{WME_{Test}} (b, \text{PositiveCondition}^{\sim}(pc), w))$$

$$\forall b : \text{Binding}; nc : \text{ran NegativeCondition}; WM : \mathcal{P} \text{WME} \bullet$$

$$\text{Matches}_{WM}^{Condition} (b, nc, WM) \Leftrightarrow$$

$$\neg (\exists nb : \text{Binding}$$

$$nb \text{ ConsistentExtension } b \wedge$$

$$nb \text{ Covers } (\text{ConditionsComponents}(nc) \cap \text{Variable}) \bullet$$

$$(\exists w : \text{WME} \bullet$$

$$\text{Matches}_{WME}^{WME_{Test}} (nb, \text{NegativeCondition}^{\sim}(nc), w)))$$

$$\forall b : \text{Binding}; ncc : \text{ran NegativeConjunctiveCondition}; WM : \mathcal{P} \text{WME} \bullet$$

$$\text{Matches}_{WM}^{Condition} (b, ncc, WM) \Leftrightarrow$$

$$\neg (\exists nb : \text{Binding}$$

$$nb \text{ ConsistentExtension } b \wedge$$

$$nb \text{ Covers } (\text{ConditionsComponents}(ncc) \cap \text{Variable}) \bullet$$

$$\forall c : \text{NegativeConjunctiveCondition}^{\sim}(ncc) \bullet$$

$$\text{Matches}_{WM}^{Condition} (nb, c, WM))$$

A match is satisfied by working memory if and only if all of its conditions match against working memory under the match's binding, and the elements they match have the current working memory numbering, i.e., the elements are the copies that are currently in memory.

$$\begin{array}{l}
 \text{Matches}_{WM}^{Match} _ : \mathcal{P}(\text{Match} \times \mathcal{P} \text{WME} \times (\text{WME} \leftrightarrow \mathbb{N})) \\
 \forall m : \text{Match}; \text{WM} : \mathcal{P} \text{WME} \bullet \\
 \text{Matches}_{WM}^{Match} (m, \text{WM}, \text{WM}\#) \Leftrightarrow \\
 (\forall c : m.\text{production.lhs} \bullet \\
 \text{Matches}_{WM}^{Condition} (m.\text{binding}, c, \text{WM}) \wedge \\
 \text{WM}\#(m.\text{matching}(c)) = \text{second}(m.\text{matching}(c)))
 \end{array}$$

5.3 Instantiation

When Soar finds a match for a production, it instantiates it and adds it to instantiation memory. The act of instantiation generates an instantiated version of the match's production's conditions and a set of preferences from the match's production's right hand side maxes. Soar adds the support conferred by the instantiation to each preference's entry in support memory. Soar checks the support of preferences to decide how to change preference memory. Soar then passes the preference memory changes to decide, which changes the contents of working memory to reflect the new contents of preference memory.

This section defines how a match's production's conditions are instantiated, and then how a match's production's right hand side is instantiated, and then the instantiation data structure schema.

5.3.1 Instantiating a Match's Production's Conditions

This section defines, in bottom-up order, seven recursive functions that instantiate the conditions of a match's production using the match's binding. The end result is a partial function, called a matching, that maps the conditions of the production to their corresponding instantiated condition.

InstantiateVariableOrConstant returns a variable's binding, if it is bound, or returns the variable, if it is unbound, and otherwise returns a constant.

$$\begin{array}{l}
 \text{InstantiateVariableOrConstant} : \\
 (\text{Variable} \cup \text{Constant}) \times \text{Binding} \rightarrow \text{Symbol} \\
 \forall cv : \text{Constant} \cup \text{Variable}; b : \text{Binding} \bullet \\
 \text{InstantiateVariableOrConstant}(cv, b) = ((\text{id } \Sigma) \oplus b)(cv)
 \end{array}$$

InstantiateSimpleTest returns tests with no arguments unmodified. Equality, same type and not tests, it destructures, instantiates their variable or constant and returns.

InstantiateSimpleTest : *SimpleTest* × *Binding* → *SimpleTest*

$\forall t : \text{SimpleTest}; b : \text{Binding} \bullet$
 $((t = \text{BlankTest} \vee t = \text{YesTest} \vee t = \text{NoTest} \vee t \in \text{ran DisjunctiveTest})$
 $\Rightarrow \text{InstantiateSimpleTest}(t, b) = t) \wedge$
 $(t \in \text{ran EqualityTest} \Rightarrow$
 $\text{InstantiateSimpleTest}(t, b) =$
 $\text{EqualityTest}(\text{InstantiateVariableOrConstant}(\text{EqualityTest}^{\sim}(t), b))) \wedge$
 $(t \in \text{ran SameTypeTest} \Rightarrow$
 $\text{InstantiateSimpleTest}(t, b) =$
 $\text{SameTypeTest}(\text{InstantiateVariableOrConstant}(\text{SameTypeTest}^{\sim}(t), b))) \wedge$
 $(t \in \text{ran NotTest} \Rightarrow$
 $\text{InstantiateSimpleTest}(t, b) =$
 $\text{NotTest}(\text{InstantiateVariableOrConstant}(\text{NotTest}^{\sim}(t), b)))$

InstantiateSetOfSimpleTests iterates the simple test instantiator across a set of tests.

InstantiateSetOfSimpleTests :
 $\mathcal{P} \text{ SimpleTest} \times \text{Binding} \rightarrow \mathcal{P} \text{ SimpleTest}$

$\forall S : \mathcal{P} \text{ SimpleTest}; b : \text{Binding} \bullet$
 $((S = \emptyset \Rightarrow \text{InstantiateSetOfSimpleTests}(S, b) = \emptyset) \wedge$
 $(S \neq \emptyset \Rightarrow$
 $(\exists s : S \bullet$
 $\text{InstantiateSetOfSimpleTests}(S, b) =$
 $\{\text{InstantiateSimpleTest}(s, b)\} \cup$
 $\text{InstantiateSetOfSimpleTests}(S \setminus \{s\}, b)))$

InstantiateTest instantiates a test by calling either *InstantiateSimpleTest* or by unpacking a conjunctive test, instantiating its component tests and re-packing.

InstantiateTest : *Test* × *Binding* → *Test*

$\forall t : \text{Test}; b : \text{Binding} \bullet$
 $((t \notin \text{ran } \text{ConjunctiveTest} \Rightarrow$
 $\quad \text{InstantiateTest}(t, b) =$
 $\quad \text{InstantiateSimpleTest}(t, b)) \wedge$
 $(t \in \text{ran } \text{ConjunctiveTest} \Rightarrow$
 $\quad \text{InstantiateTest}(t, b) =$
 $\quad \text{ConjunctiveTest}(\text{InstantiateSetOfSimpleTests}(\text{ConjunctiveTest}^{\sim}(t), b)))$

InstantiateWMETest instantiates a working memory element test by instantiating its component parts, and giving the new test the same context acceptable preference test.

InstantiateWMETest : *WMETest* × *Binding* → *WMETest*

$\forall w : \text{WMETest}; b : \text{Binding} \bullet$
 $\text{InstantiateWMETest}(w, b) =$
 $(\mu \text{id} : \{\text{InstantiateTest}(w.\text{id}, b)\} \bullet$
 $(\mu \text{attribute} : \{\text{InstantiateTest}(w.\text{attribute}, b)\} \bullet$
 $(\mu \text{value} : \{\text{InstantiateTest}(w.\text{value}, b)\} \bullet$
 $(\mu \text{WMETest} :$
 $\quad \text{context_acceptable_preference} =$
 $\quad w.\text{context_acceptable_preference} \bullet$
 $\quad \theta \text{WMETest}))))$

InstantiateCondition and *InstantiateSetOfConditions* recurse through conditions instantiating their component working memory element tests or conditions. Z requires that they are defined in a single axiomatic definition to allow them to be mutually recursive.

$$\begin{array}{l}
 \textit{InstantiateCondition} : \textit{Condition} \times \textit{Binding} \rightarrow \textit{Condition} \\
 \textit{InstantiateSetOfConditions} : \mathbb{P} \textit{Condition} \times \textit{Binding} \rightarrow \mathbb{P} \textit{Condition} \\
 \hline
 \forall c : \textit{Condition}; b : \textit{Binding} \bullet \\
 ((c \in \textit{ran PositiveCondition} \Rightarrow \\
 \textit{InstantiateCondition}(c, b) = \\
 \textit{PositiveCondition}(\textit{InstantiateWMETest}(\textit{PositiveCondition}^{\sim}(c), b))) \wedge \\
 (c \in \textit{ran NegativeCondition} \Rightarrow \\
 \textit{InstantiateCondition}(c, b) = \\
 \textit{NegativeCondition}(\textit{InstantiateWMETest}(\textit{NegativeCondition}^{\sim}(c), b))) \wedge \\
 (c \in \textit{ran NegativeConjunctiveCondition} \Rightarrow \\
 \textit{InstantiateCondition}(c, b) = \\
 \textit{NegativeConjunctiveCondition} \\
 (\textit{InstantiateSetOfConditions} \\
 (\textit{NegativeConjunctiveCondition}^{\sim}(c), b)))) \\
 \forall S : \mathbb{P} \textit{Condition}; b : \textit{Binding} \bullet \\
 ((S = \emptyset \Rightarrow \textit{InstantiateSetOfConditions}(S, b) = \emptyset) \wedge \\
 (S \neq \emptyset \Rightarrow \\
 (\exists s : S \bullet \\
 \textit{InstantiateSetOfConditions}(S, b) = \\
 \{\textit{InstantiateCondition}(s, b)\} \cup \\
 \textit{InstantiateSetOfConditions}(S \setminus \{s\}, b))))
 \end{array}$$

InstantiateConditionsWME instantiates a condition that corresponds to the working memory element that a positive condition matched.

InstantiateConditionsWME : *WME* \rightarrow *Condition*

$\forall w : WME \bullet$

$((w.context_acceptable_preference = Yes \Rightarrow$

InstantiateConditionsWME(*w*) =

PositiveCondition(

makeWMETest(*EqualityTest*(*w.id*), *EqualityTest*(*w.attribute*),

EqualityTest(*w.value*), *YesTest*)))

\wedge

$(w.context_acceptable_preference = Yes \Rightarrow$

InstantiateConditionsWME(*w*) =

PositiveCondition(

makeWMETest(*EqualityTest*(*w.id*), *EqualityTest*(*w.attribute*),

EqualityTest(*w.value*), *NoTest*)))

InstantiateConditionMatching instantiates a set of conditions under a binding to produce a partial function that maps each of the conditions to their instantiation. The positive conditions are mapped to a simple instantiated version of the working memory element that they matched. It contains an equality test for each value in the element and the appropriate context acceptable preference test. The negative conditions are recursively instantiated to retain their test structures, but have their variables replaced with identifiers. This retains maximal information about their matching which can be utilized when the chunk is built.

$$\begin{array}{l}
 \textit{InstantiateConditionMatching} : \\
 (\text{ran } \textit{PositiveCondition} \rightarrow (\textit{WME} \times \mathbb{N})) \times \\
 \mathbb{P}(\textit{Condition} \setminus \text{ran } \textit{PositiveCondition}) \times \textit{Binding} \\
 \rightarrow (\textit{Condition} \rightarrow \textit{Condition}) \\
 \hline
 \forall \textit{CM} : (\text{ran } \textit{PositiveCondition} \rightarrow (\textit{WME} \times \mathbb{N})); \textit{C} : \mathbb{P} \textit{Condition}; \textit{b} : \textit{Binding} \bullet \\
 ((\textit{CM} = \emptyset \wedge \textit{C} = \emptyset \Rightarrow \textit{InstantiateConditionMatching}(\textit{C}, \textit{b}) = \emptyset) \wedge \\
 (\textit{CM} \neq \emptyset \Rightarrow \\
 (\exists \textit{cwm} : \textit{CM} \bullet \\
 \textit{InstantiateConditionMatching}(\textit{CM}, \textit{C}, \textit{b}) = \\
 \{(\textit{first}(\textit{cwm}) \rightarrow \textit{InstantiateConditionsWME}(\textit{first}(\textit{second}(\textit{cwm}))) \} \cup \\
 \textit{InstantiateConditionMatching}(\textit{CM} \setminus \textit{cwm}, \textit{C}, \textit{b})) \wedge \\
 (\textit{CM} = \emptyset \wedge \textit{C} \neq \emptyset \Rightarrow \\
 (\exists \textit{c} : \textit{C} \bullet \\
 \textit{InstantiateConditionMatching}(\textit{CM}, \textit{C}, \textit{b}) = \\
 \{(\textit{c} \rightarrow \textit{InstantiateNegativeCondition}(\textit{c}, \textit{b}))\} \cup \\
 \textit{InstantiateConditionMatching}(\emptyset, \textit{C} \setminus \{\textit{c}\}, \textit{b}))
 \end{array}$$

5.3.2 Instantiating the RHS of a Match's Production

MakesPreference takes a binding and produces a function that translates a make into a preference using the binding. The different argument format allows the function to be easily imaged across a set. The binding must cover all of the variables of an argument preference, or *MakesPreference* is undefined.

$$\text{MakesPreference} : \text{Binding} \rightarrow (\text{Make} \rightarrow \text{Preference})$$

$$\forall b : \text{Binding} \bullet \forall m : \text{Make} \bullet$$

$$\exists o : \Sigma \rightarrow \Sigma \quad o = (\text{id Constant} \oplus b) \bullet$$

$$((m \in \text{ran UnaryM} \Rightarrow$$

$$(\exists up : \text{UnaryPreference} \mid up = \text{UnaryM}^{\sim}(m) \bullet$$

$$\text{MakesPreference}(b)(m) =$$

$$\text{makeUnaryPreference}(o(up.id), o(up.attribute),$$

$$o(up.value), up.preference))) \wedge$$

$$(m \in \text{ran BinaryM} \Rightarrow$$

$$(\exists bp : \text{BinaryPreference} \mid bp = \text{BinaryM}^{\sim}(m) \bullet$$

$$\text{MakesPreference}(b)(m) =$$

$$\text{makeBinaryPreference}(o(bp.id), o(bp.attribute), o(bp.value),$$

$$bp.preference, o(bp.referent))))))$$

5.3.3 Instantiation

An instantiation contains a match, a binding and a matching. The binding is a consistent extension of the match's binding that covers all of the variables of the production's right hand side. The instantiation's right hand side is the set of preferences made from instantiating the production's makes under the instantiation's binding.

The instantiation also contains a matching. The matching is a correspondence between *all* of the conditions of the production and new conditions. The new conditions are instantiated copies of the production's conditions under the match's binding; these record for chunking how the conditions of the productions matched.

<p><i>Instantiation</i></p> <p><i>match</i> : <i>Match</i></p> <p><i>binding</i> : <i>Binding</i></p> <p><i>matching</i> : <i>Condition</i> \rightarrow <i>Condition</i></p> <p><i>lhs</i> : \mathbb{P} <i>Condition</i></p> <p><i>rhs</i> : \mathbb{P} <i>Preference</i></p> <hr/> <p><i>binding</i> = <i>ConsistentExtension match.binding</i></p> <p><i>dom binding</i> = $(\bigcup (\text{MakesComponents}(\text{match.production.rhs})) \cap \text{Variable})$ $\cup \text{dom match.binding}$</p> <p><i>rhs</i> = $(\text{MakesPreference}(\text{binding}))(\text{match.production.rhs})$</p> <p><i>matching</i> = <i>InstantiateConditionMatching</i>(<i>match.matching</i>, $\text{match.production.lhs} \setminus \text{ran PositiveCondition}$, <i>match.binding</i>)</p> <p><i>lhs</i> = ran matching</p>

makeInstantiation constructs a new instantiation data structure schema given a match and a condition matching. The instantiation's binding, left hand side and right hand side are all uniquely determined from the match and the condition matching.

<p><i>makeInstantiation</i> : <i>Match</i> \times (<i>Condition</i> \rightarrow <i>Condition</i>) \rightarrow <i>Instantiation</i></p> <hr/> <p>\forall <i>match</i> : <i>Match</i>; <i>matching</i> : (<i>Condition</i> \rightarrow <i>Condition</i>) • \exists_1 <i>binding</i> : <i>Binding</i>; <i>lhs</i> : \mathbb{P} <i>Condition</i>; <i>rhs</i> : \mathbb{P} <i>Preference</i> <i>Instantiation</i> • <i>makeInstantiation</i>(<i>match</i>, <i>matching</i>) = θ<i>Instantiation</i></p>
--

5.4 Goal Operations and a Match's Goal

Recognition memory reads the contents of goal memory and working memory to determine the support that instantiations confer on preferences. The goal stack is defined in Chapter 7, but is represented in working memory elements. This section defines operations that read from working memory the state, operator and ancestors of a goal, and determine the newest goal that a match matches against.

The state of a goal is the value of the single working memory element of the goal that has attribute state.

$$\text{GoalState} : \text{Identifier} \times \mathbb{P} \text{ TME} \rightarrow \text{Identifier}$$

$$\begin{array}{l} \forall g : \text{Identifier}; TM : \mathbb{P} \text{ TME} \bullet \\ \exists_1 t : TM \cap \text{ran } \text{WMETME} \bullet \\ \quad \exists w : \text{WME} \mid w = \text{WMETME}^{\sim}(t) \wedge \\ \quad \quad w.\text{id} = g \wedge w.\text{attribute} = \text{'STATE'} \wedge \\ \quad \quad w.\text{context_acceptable_preference} = \text{No} \\ \bullet \text{GoalState}(g, TM) = w.\text{value} \end{array}$$

The operator of a goal is the value of the single working memory element of the goal that has attribute operator.

$$\text{GoalOperator} : \text{Identifier} \times \mathbb{P} \text{ TME} \rightarrow \text{Identifier}$$

$$\begin{array}{l} \forall g : \text{Identifier}; TM : \mathbb{P} \text{ TME} \bullet \\ \exists t : TM \cap \text{ran } \text{WMETME} \bullet \\ \quad \exists w : \text{WME} \mid w = \text{WMETME}^{\sim}(t) \wedge \\ \quad \quad w.\text{id} = g \wedge w.\text{attribute} = \text{'OPERATOR'} \wedge \\ \quad \quad w.\text{context_acceptable_preference} = \text{No} \\ \bullet \text{GoalOperator}(g, TM) = w.\text{value} \end{array}$$

GoalAncestors returns the set of identifiers of ancestor goals of a goal.

$$\text{GoalAncestors} : \text{Identifier} \times \mathbb{P} \text{ TME} \rightarrow \mathbb{P} \text{ Identifier}$$

$$\begin{array}{l} \forall g : \text{Identifier}; TM : \mathbb{P} \text{ TME} \bullet \\ \text{GoalAncestors}(\text{'NIL'}, TM) = \emptyset \wedge \\ (\exists g_2 : \text{Identifier} \mid \text{WMETME}(\text{makeWME}(g, \text{'OBJECT'}, g_2, \text{No})) \in TM \bullet \\ \text{GoalAncestors}(g, TM) = \{g_2\} \cup \text{GoalAncestors}(g_2, TM)) \end{array}$$

The support an instantiation confers depends upon the deepest of the goals that it matched, called the match's goal. Soar constrains its productions to match each root variable with an identifier of a goal or attribute impasse. A root variable is one that appears only in the identifier field of the production's conditions.

$MatchGoals : Match \rightarrow \mathbb{P} Identifier$ $MatchGoal : Match \times \mathbb{P} TME \rightarrow Identifier$
$\forall m : Match \bullet$ $MatchGoals(m) = m.binding(\{m.production.roots\})$
$\forall m : Match; TM : \mathbb{P} TME \bullet$ $\exists g : MatchGoals(m) \bullet$ $\neg (\exists g_2 : MatchGoals(m) \bullet g_2 \neq g \wedge g \in GoalAncestors(g_2, TM)) \bullet$ $MatchGoal(m, TME) = g$

5.5 Chunking

Chunking is Soar's explanation based learning mechanism [RL86]. It watches and records some of recognition memory's and decide's operation into two trace memories. When decide fills a slot in working memory, it augments the working memory trace, *TrW*. *TrW* maps each working memory element instance to the require or acceptable preference for the slot and value. After recognition memory instantiates a match, it passes the instantiation to chunking. Chunking creates a data structure called a trace from the instantiation. The trace records the working memory element instances that were matched in the form of instantiated conditions. The instantiated conditions are split into three groups, called grounds, potentials and locals. Each preference instance of the instantiation is marked as dependent upon this trace in the preference memory trace, *TrP*.

After chunking has created the instantiation's trace, it determines if the instantiation augments the match goal's parent with new preferences. If so, the results are used as a starting point for a search through the trace memories. The trace generated conditions in this search are collected and variabilized to produce a new left hand side for a production, and the results are variabilized to produce a new right hand side.

Chunking is presented in eleven sections.

- Section 5.5.1 Extracting Not Tests — instantiations test values from their working memory elements for inequality. The resulting chunks must also check these inequalities, so we extract the not tests of each instantiation, store them in the traces and re-introduce them in Section 5.5.3 Not-ify.
- Section 5.5.2 Variabilize — this section generalizes the sets of instantiated

conditions that chunking finds into conditions by assigning variables to each identifier.

- Section 5.5.3 Not-ify — re-introduces not tests.
- Section 5.5.4 Traces — specifies what is in a trace and provides utilities for constructing them.
- Section 5.5.5 Chunking's State — this section defines chunking's state schema.
- Section 5.5.6 Changing, Initializing and Resetting Chunking — we define constraints on each step of chunking, as well as the standard initialization and resetting for chunking.
- Section 5.5.7 Defining Results — results are the preferences that an instantiation adds to its parent goal, which are also the starting point for chunking's backtracing.
- Section 5.5.8 Chunking an Instantiation — this section provides the operation to build the trace for an instantiation, and to start the search back through the trace memories.
- Section 5.5.9 Tracing Seeds, Locals and Grounded Potentials — this section specifies the basic tracing operations.
- Section 5.5.10 Tracing UnGrounded Potentials — this section defines how to trace a special type of condition that may or may not be included in the chunk's condition.
- Section 5.5.11 ChunkSP — this section specifies how chunking builds the new chunk when backtracing has finished.

5.5.1 Extracting Not Tests

In this section, we specify how to extract the nots from a condition. The nots are doubletons of identifiers that the condition constrained to be different. The presentation is bottom up, starting with a not extractor for tests, then working memory element tests and then condition's nots.

$$Not == \{c_1, c_2 : Identifier \mid c_1 \neq c_2 \bullet \{c_1, c_2\}\}$$

Most types of tests have no nots, but a *NotTest* that matched against two identifiers generates a not doubleton.

$$\text{TestsNots} : \text{Binding} \times \text{Test} \times \text{Test} \rightarrow \mathbb{P} \text{Not}$$

$$\begin{aligned} &\forall b : \text{Binding}; t, vt : \text{Test} \bullet \\ &((t \in \{\text{BlankTest}, \text{YesTest}, \text{NoTest}\} \Rightarrow \text{TestsNots}(b, t, vt) = \emptyset) \wedge \\ &(t \in \text{ran EqualityTest} \Rightarrow \text{TestsNots}(b, t, vt) = \emptyset) \wedge \\ &(t \in \text{ran NotTest} \Rightarrow \\ &(\exists s : \Sigma \mid s = \text{NotTest} \sim (t) \bullet \\ &(((s \in \text{Variable}) \wedge (((\text{id Constant}) \oplus b)(s) \in \text{Identifier}) \\ &\wedge (\text{EqualityTest} \sim (vt) \in \text{Identifier})) \Rightarrow \\ &\text{TestsNots}(b, t, vt) = \{\{((\text{id Constant}) \oplus b)(s), \text{EqualityTest} \sim (vt)\}\}) \wedge \\ &(((s \in \text{Variable}) \wedge (((\text{id Constant}) \oplus b)(s) \in \text{Identifier}) \\ &\wedge (\text{EqualityTest} \sim (vt) \in \text{Identifier})) \Rightarrow \\ &\text{TestsNots}(b, t, vt) = \emptyset))) \wedge \\ &(t \in \text{ran SameTypeTest} \Rightarrow \text{TestsNots}(b, t, vt) = \emptyset) \wedge \\ &(t \in \text{ran DisjunctiveTest} \Rightarrow \text{TestsNots}(b, t, vt) = \emptyset) \wedge \\ &(t \in \text{ran ConjunctiveTest} \Rightarrow \\ &\text{TestsNots}(b, t, vt) = \\ &\bigcup((\lambda t : \text{Test} \bullet \text{TestsNots}(b, t, vt))(\text{ConjunctiveTest} \sim (t)))) \end{aligned}$$

The nots of a working memory element test is the set of nots in its id, attribute and value tests.

$$\text{WMETestsNots} : \text{Binding} \times \text{WMETest} \times \text{WMETest} \rightarrow \mathbb{P} \text{Not}$$

$$\begin{aligned} &\forall b : \text{Binding}; wt, w : \text{WMETest} \bullet \\ &\text{WMETestsNots}(b, wt, w) = \\ &\text{TestsNots}(b, wt.\text{id}, w.\text{id}) \cup \text{TestsNots}(b, wt.\text{attribute}, w.\text{attribute}) \\ &\cup \text{TestsNots}(b, wt.\text{value}, w.\text{value}) \end{aligned}$$

We only extract nots from positive conditions, as the nots inside of a negative condition are known not to have matched against any values. The nots of a positive condition are the nots of its working memory element tests.

$$\text{ConditionsNots} : \text{Binding} \times \text{Condition} \times \text{Condition} \rightarrow \mathbb{P} \text{Not}$$

$$\begin{aligned} &\forall b : \text{Binding}; c, cm : \text{Condition} \bullet \\ &((c \in \text{ran PositiveCondition} \Rightarrow \\ &\text{ConditionsNots}(b, c, cm) = \\ &\text{WMETestsNots}(b, \text{PositiveCondition} \sim (c), \\ &\text{PositiveCondition} \sim (cm))) \wedge \\ &(c \notin \text{ran PositiveCondition} \Rightarrow \\ &\text{ConditionsNots}(b, c, cm) = \emptyset)) \end{aligned}$$

5.5.2 Variabilize

Variabilize generalizes Chunking's instantiated conditions by replacing their identifiers with variables. After chunking has finished its search through the trace memory, it picks an assignment, a partial injection from identifiers to variables, that covers all of the variables of its condition. It recurses down into the tests of its conditions, replacing all occurrences of identifiers with their variables.

$$\text{Assignment} == \text{Identifier} \mapsto \text{Variable}$$

A symbol is variabilized by returning the symbol if it is a constant and otherwise by returning its value under the assignment.

$$\text{VariabilizeSymbol} : \Sigma \times \text{Assignment} \rightarrow \Sigma$$

$$\begin{aligned} &\forall c : \Sigma; a : \text{Assignment} \bullet \\ &((c \in \text{Identifier} \Rightarrow \\ &(\exists v : \text{Variable} \mid v \notin \text{ran } a \bullet \text{VariabilizeSymbol}(c, a) = v)) \wedge \\ &(c \notin \text{Identifier} \Rightarrow \text{VariabilizeSymbol}(c, a) = c)) \end{aligned}$$

A test is variabilized by variabilizing its component symbols, or returning the test if it is indivisible.

$$\text{VariabilizeTest} : \text{Test} \times \text{Assignment} \rightarrow \text{Test}$$

$$\begin{aligned} &\forall t : \text{Test}; a : \text{Assignment} \bullet \\ &((t \in \{\text{BlankTest}, \text{YesTest}, \text{NoTest}\} \cup \text{ran DisjunctiveTest} \\ &\Rightarrow \text{VariabilizeTest}(t, a) = t) \wedge \\ &(t \in \text{ran EqualityTest} \Rightarrow \\ &\text{VariabilizeTest}(t, a) = \\ &\text{EqualityTest}(\text{VariabilizeSymbol}(\text{EqualityTest}^{\sim}(t), a))) \wedge \\ &(t \in \text{ran NotTest} \Rightarrow \\ &\text{VariabilizeTest}(t, a) = \\ &\text{NotTest}(\text{VariabilizeSymbol}(\text{NotTest}^{\sim}(t), a))) \wedge \\ &(t \in \text{ran SameTypeTest} \Rightarrow \\ &\text{VariabilizeTest}(t, a) = \\ &\text{SameTypeTest}(\text{VariabilizeSymbol}(\text{SameTypeTest}^{\sim}(t), a))) \wedge \\ &(t \in \text{ran ConjunctiveTest} \Rightarrow \\ &\text{VariabilizeTest}(t, a) = \\ &\text{ConjunctiveTest}(((\lambda st : \text{Test} \bullet \text{VariabilizeTest}(st, a)) \\ &\quad \Downarrow \text{ConjunctiveTest}^{\sim}(t)))))) \end{aligned}$$

A working memory element test is variabilized by variabilizing its id, attribute and value tests, and keeping its acceptable preference test.

$$\text{Variabilize WMETest} : \text{WMETest} \times \text{Assignment} \rightarrow \text{WMETest}$$

$$\begin{aligned} \forall w : \text{WMETest}; a : \text{Assignment} \bullet \\ \text{Variabilize WMETest}(w, a) = \\ (\mu \text{id} : \{ \text{Variabilize Test}(w.\text{id}, a) \} \bullet \\ (\mu \text{attribute} : \{ \text{Variabilize Test}(w.\text{attribute}, a) \} \bullet \\ (\mu \text{value} : \{ \text{Variabilize Test}(w.\text{value}, a) \} \bullet \\ (\mu \text{context_acceptable_preference} : \{ w.\text{context_acceptable_preference} \} \bullet \\ \theta \text{ WMETest})))))) \end{aligned}$$

A condition is variabilized by variabilizing its recursive set of conditions or its component working memory element test.

$$\text{Variabilize Condition} : \text{Condition} \times \text{Assignment} \rightarrow \text{Condition}$$

$$\text{Variabilize SetOfConditions} :$$

$$\mathbb{P} \text{Condition} \times \text{Assignment} \rightarrow \mathbb{P} \text{Condition}$$

$$\begin{aligned} \forall c : \text{Condition}; a : \text{Assignment} \bullet \\ ((c \in \text{ran PositiveCondition} \Rightarrow \\ \text{Variabilize Condition}(c, a) = \\ \text{PositiveCondition}(\text{Variabilize WMETest}(\text{PositiveCondition}^{\sim}(c), a))) \wedge \\ (c \in \text{ran NegativeCondition} \Rightarrow \\ \text{Variabilize Condition}(c, a) = \\ \text{NegativeCondition}(\text{Variabilize WMETest}(\text{NegativeCondition}^{\sim}(c), a))) \wedge \\ (c \in \text{ran NegativeConjunctiveCondition} \Rightarrow \\ \text{Variabilize Condition}(c, a) = \\ \text{NegativeConjunctiveCondition}(\text{Variabilize SetOfConditions}(\text{NegativeConjunctiveCondition}^{\sim}(c), a)))) \\ \forall C : \mathbb{P} \text{Condition}; a : \text{Assignment} \bullet \\ ((C = \emptyset \Rightarrow \text{Variabilize SetOfConditions}(C, a) = \emptyset) \wedge \\ (C \neq \emptyset \Rightarrow \\ (\exists c : C \bullet \\ \text{Variabilize SetOfConditions}(C, a) = \\ \{ \text{Variabilize Condition}(c, a) \} \cup \text{Variabilize SetOfConditions}(C \setminus \{c\}, a)))))) \end{aligned}$$

Chunking cannot produce a set of instantiated conditions to variabilize because it must maintain the correspondence between the conditions and the working memory element instances that they matched. Chunking actually produces a partial function from conditions to working memory element instances. This is variabilized to produce both a matching for the chunk's match and another form for the chunk's instantiation.

$$\frac{\text{VariabilizeConditionMatching} : (\text{Condition} \rightarrow (\text{WME} \times \mathbb{N})) \times \text{Assignment} \rightarrow ((\text{Condition} \rightarrow (\text{WME} \times \mathbb{N})) \times (\text{Condition} \rightarrow \text{Condition}))}{\text{VariabilizeConditionMatching} : (\text{Condition} \rightarrow (\text{WME} \times \mathbb{N})) \times \text{Assignment} \rightarrow ((\text{Condition} \rightarrow (\text{WME} \times \mathbb{N})) \times (\text{Condition} \rightarrow \text{Condition}))}$$

$$\begin{aligned} & \forall mm : (\text{Condition} \rightarrow (\text{WME} \times \mathbb{N})); a : \text{Assignment} \bullet \\ & ((mm = \emptyset \Rightarrow \text{VariabilizeConditionMatching}(mm, a) = (\emptyset, \emptyset)) \wedge \\ & (mm \neq \emptyset \Rightarrow \\ & (\exists c : \text{dom } mm \bullet \\ & (\exists cwncc : \{ \text{VariabilizeConditionMatching}(\{c\} \ll mm, a) \} \bullet \\ & \text{VariabilizeConditionMatching}(mm, a) = \\ & (\{ \text{VariabilizeCondition}(c, a) \rightarrow mm(c) \} \cup \text{first}(cwncc), \\ & \{ \text{VariabilizeCondition}(c, a) \rightarrow c \} \cup \text{second}(cwncc)))))) \end{aligned}$$

The results of an instantiation are variabilized to produce the new right hand side for the chunk. *VariabilizePreferenceToMake* maps a preference to a make. and *VariabilizeRHS* maps a set of preferences to its set of makes.

VariabilizePreferenceToMake : *Preference* × *Assignment* → *Make*

$$\begin{aligned} &\forall p : \text{Preference}; a : \text{Assignment} \bullet \\ &((p \in \text{ran UnaryP} \Rightarrow \\ &(\exists up : \text{UnaryPreference} \mid up = \text{UnaryP}^{\sim}(p) \bullet \\ &\text{VariabilizePreferenceToMake}(p, a) = \\ &\text{makeUnaryMake}(\text{VariabilizeSymbol}(up.id, a), \\ &\text{VariabilizeSymbol}(up.attribute, a), \\ &\text{VariabilizeSymbol}(up.value, a), up.preference))) \wedge \\ &(p \in \text{ran BinaryP} \Rightarrow \\ &(\exists bp : \text{BinaryPreference} \mid bp = \text{BinaryP}^{\sim}(p) \bullet \\ &\text{VariabilizePreferenceToMake}(p, a) = \\ &\text{makeBinaryMake}(\text{VariabilizeSymbol}(bp.id, a), \\ &\text{VariabilizeSymbol}(bp.attribute, a), \\ &\text{VariabilizeSymbol}(bp.value, a), bp.preference, \\ &\text{VariabilizeSymbol}(bp.referent, a)))))) \end{aligned}$$

VariabilizeRHS : \mathbb{P} *Preference* × *Assignment* → \mathbb{P} *Make*

$$\begin{aligned} &\forall ps : \mathbb{P} \text{ Preference}; a : \text{Assignment} \bullet \\ &((ps = \emptyset \Rightarrow \text{VariabilizeRHS}(ps, a) = \emptyset) \wedge \\ &(ps \neq \emptyset \Rightarrow \\ &(\exists p : ps \bullet \\ &\text{VariabilizeRHS}(ps, a) = \\ &\{\text{VariabilizePreferenceToMake}(p, a)\} \cup \\ &\text{VariabilizeRHS}(ps \setminus \{p\}, a)))) \end{aligned}$$

5.5.3 Not-ify

When a trace for an instantiation is created, the not extractor saves the set of doubletons of identifiers that the match's condition constrained to be different. As chunking searches through the memory of traces, it unions in the set of not doubletons for each trace that it adds. When chunking is completed, not-ify re-installs these nots into the new conditions.

Not-ify is called after variabilize, so we map the set of doubletons of identifiers to a set of doubletons of variables, called *VNot*s. The function *NotstoVNot*s maps set of identifier nots to variable nots using the chunk's assignment.

$$VNot == \{v_1, v_2 : Variable \mid v_1 \neq v_2 \bullet \{v_1, v_2\}\}$$

$$NotstoVNot : \mathcal{P} Not \times Assignment \rightarrow \mathcal{P} VNot$$

$$\forall N : \mathcal{P}_1 Not; a : Assignment \bullet$$

$$NotstoVNot(N, a) =$$

$$\{n : N; i, j : Identifier \mid i \neq j \wedge n = \{i, j\} \bullet \{a(i), a(j)\}\}$$

Not-ify's functions all take three arguments: an object to not-ify, the set of variable nots to install and the set of variables bound in the walk through the conditions. As not-ify walks down the condition matchings, into the conditions, working memory element tests and tests, it adds the variables that the conditions have bound to the set of variables. As it descends, it also installs not tests for those variable nots that have had both of their variables bound. Each not-ify function returns the possibly augmented set of bound variables, and the possibly reduced set of variable nots left to be installed.

As the domains and ranges of the not-ify functions must conveniently handle triples of items, we extend the standard *Z first* and *second* to sequences of three elements ([Spi89] 93).

$$= [X, Y, Z]$$

$$first_3 : X \times Y \times Z \rightarrow X$$

$$second_3 : X \times Y \times Z \rightarrow Y$$

$$third_3 : X \times Y \times Z \rightarrow Z$$

$$\forall x : X; y : Y; z : Z \bullet first_3((x, y, z)) = x$$

$$\forall x : X; y : Y; z : Z \bullet second_3((x, y, z)) = y$$

$$\forall x : X; y : Y; z : Z \bullet third_3((x, y, z)) = z$$

NotifyTest returns most tests unmodified, but if a test is an equality for an unbound variable it binds the variable by adding it to the variable set, and recurses to possibly add in an as of yet unassigned variable not. *NotifySetOfTests* sequences through a set of tests in non-deterministic order to install the variable nots.

NotifyTest :

$Test \times \mathbb{P} VNot \times \mathbb{P} Variable \rightarrow$

$Test \times \mathbb{P} VNot \times \mathbb{P} Variable$

NotifySetOfTests :

$\mathbb{P} Test \times \mathbb{P} VNot \times \mathbb{P} Variable \rightarrow$

$\mathbb{P} Test \times \mathbb{P} VNot \times \mathbb{P} Variable$

$\forall t : Test; N : \mathbb{P} VNot; V : \mathbb{P} Variable \bullet$

$((t \in \text{ran } EqualityTest \Rightarrow$

$(EqualityTest \sim (t) \notin Variable \Rightarrow$

$NotifyTest(t, N, V) = (t, N, V)) \wedge$

$(EqualityTest \sim (t) \in Variable \Rightarrow$

$(\exists v : Variable \mid v = EqualityTest \sim (t) \bullet$

$(\exists n : N \bullet$

$v \in n \wedge (n \setminus \{v\}) \cap V \neq \emptyset \Rightarrow$

$(\exists n : N; v_2 : Variable \mid n = \{v, v_2\} \wedge v_2 \in V \bullet$

$(\exists r : \{NotifyTest(ConjunctiveTest(\{t, NotTest(v_2)\}), N \setminus \{n\}, V)\} \bullet$

$NotifyTest(t, N, V) = NotifyTest(first_3(r), second_3(r), third_3(r)))) \wedge$

$\neg (\exists n : N \bullet v \in n \wedge (n \setminus \{v\}) \cap V \neq \emptyset \Rightarrow$

$NotifyTest(t, N, V) = (t, N, V \cup \{v\})))) \wedge$

$(t \in \text{ran } ConjunctiveTest \Rightarrow$

$(\exists tNV : \{NotifySetOfTests(ConjunctiveTest \sim (t), N, V)\} \bullet$

$NotifyTest(t, N, V) = (ConjunctiveTest(first_3(tNV), second_3(tNV), third_3(tNV)))) \wedge$

$(t \notin (\text{ran } EqualityTest \cup \text{ran } ConjunctiveTest) \Rightarrow$

$NotifyTest(t, N, V) = (t, N, V)))$

$\forall T : \mathbb{P} Test; N : \mathbb{P} VNot; V : \mathbb{P} Variable \bullet$

$((T = \emptyset \Rightarrow NotifySetOfTests(T, N, V) = (T, N, V)) \wedge$

$(T \neq \emptyset \Rightarrow$

$(\exists t : T \bullet$

$(\exists tNV : \{NotifyTest(t, N, V)\} \bullet$

$(\exists TNV : \{NotifySetOfTests(T \setminus \{t\}, second_3(tNV), third_3(tNV))\} \bullet$

$NotifySetOfTests(T, N, V) =$

$(\{t\} \cup first_3(TNV), second_3(TNV), third_3(TNV))))))$

NotifyWMETest sequences through the id, attribute and value to install the variable nots.

$$\begin{array}{l} \text{NotifyWMETest} : \text{WMETest} \times \mathbb{P} \text{VNot} \times \mathbb{P} \text{Variable} \rightarrow \\ \text{WMETest} \times \mathbb{P} \text{VNot} \times \mathbb{P} \text{Variable} \end{array}$$

$$\begin{array}{l} \forall w : \text{WMETest}; n : \mathbb{P} \text{VNot}; b : \mathbb{P} \text{Variable} \bullet \\ \text{NotifyWMETest}(w, n, b) = \\ (\mu \text{idnb} : \{\text{NotifyTest}(w.\text{id}, n, b)\} \bullet \\ (\mu \text{attributenb} : \{\text{NotifyTest}(w.\text{attribute}, \text{second}_3(\text{idnb}), \text{third}_3(\text{idnb}))\} \bullet \\ (\mu \text{attribute} : \{\text{first}_3(\text{attributenb})\} \bullet \\ (\mu \text{valuenb} : \{\text{NotifyTest}(w.\text{value}, \text{second}_3(\text{attributenb}), \text{third}_3(\text{attributenb}))\} \bullet \\ (\text{makeWMETest}(\text{first}_3(\text{idnb}), \text{first}_3(\text{attributenb}), \\ \text{first}_3(\text{valuenb}), w.\text{context_acceptable_preference}, \\ \text{second}_3(\text{valuenb}), \text{third}_3(\text{valuenb})))))) \end{array}$$

NotifyCondition unpacks conditions positive conditions to install variable nots, but does not install nots in negated conditions.

$$\begin{array}{l} \text{NotifyCondition} : \\ \text{Condition} \times \mathbb{P} \text{VNot} \times \mathbb{P} \text{Variable} \rightarrow \text{Condition} \times \mathbb{P} \text{VNot} \times \mathbb{P} \text{Variable} \end{array}$$

$$\begin{array}{l} \forall c : \text{Condition}; n : \mathbb{P} \text{VNot}; b : \mathbb{P} \text{Variable} \bullet \\ ((c \in \text{ran PositiveCondition} \Rightarrow \\ (\exists \text{pcnb} : \{\text{NotifyWMETest}(\text{PositiveCondition} \sim (c), n, b)\} \bullet \\ \text{NotifyCondition}(c, n, b) = \\ (\text{PositiveCondition}(\text{first}_3(\text{pcnb}), \\ \text{second}_3(\text{pcnb}), \text{third}_3(\text{pcnb})))))) \wedge \\ (c \notin \text{ran PositiveCondition} \Rightarrow \\ \text{NotifyCondition}(c, n, b) = (c, n, b)) \end{array}$$

Finally, *NotifyConditionMatchings* recurses through the two style of condition matchings to install all of the variable nots.

$$\begin{array}{l}
 \text{NotifyConditionMatchings :} \\
 ((\text{Condition} \rightarrow (\text{WME} \times \mathbb{N})) \times (\text{Condition} \rightarrow \text{Condition})) \times \mathbb{P} \text{VNot} \times \mathbb{P} \text{Variable} \rightarrow \\
 ((\text{Condition} \rightarrow (\text{WME} \times \mathbb{N})) \times (\text{Condition} \rightarrow \text{Condition})) \times \mathbb{P} \text{VNot} \times \mathbb{P} \text{Variable} \\
 \hline
 \forall \text{cwn} : (\text{Condition} \rightarrow (\text{WME} \times \mathbb{N})); \text{cc} : (\text{Condition} \rightarrow \text{Condition}); \\
 \text{N} : \mathbb{P} \text{VNot}; \text{V} : \mathbb{P} \text{Variable} \bullet \\
 ((\text{cwn} = \emptyset \wedge \text{cc} = \emptyset \Rightarrow \\
 \text{NotifyConditionMatchings}((\text{cwn}, \text{cc}), \text{N}, \text{V}) = ((\emptyset, \emptyset), \text{N}, \text{V})) \wedge \\
 (\text{cwn} \neq \emptyset \wedge \text{cc} \neq \emptyset \Rightarrow \\
 (\exists c : \text{dom cwn} \cap \text{dom cc} \bullet \\
 (\exists \text{vcNV} : \{\text{NotifyCondition}(c, \text{N}, \text{V})\} \bullet \\
 (\exists \text{cnccNV} : \{\text{NotifyConditionMatchings} \\
 (\{c\} \Leftarrow \text{cwn}, \{c\} \Leftarrow \text{cc}), \text{second}_3(\text{vcNV}), \text{third}_3(\text{vcNV})\}) \bullet \\
 \text{NotifyConditionMatchings}((\text{cwn}, \text{cc}), \text{N}, \text{V}) = \\
 ((\text{first}(\text{first}_3(\text{cnccNV})) \oplus \{\text{first}_3(\text{vcNV}) \mapsto \text{cwn}(c)\}, \\
 \text{second}(\text{first}_3(\text{cnccNV})) \oplus \{\text{first}_3(\text{vcNV}) \mapsto c\}), \\
 \text{second}_3(\text{cnccNV}), \text{third}_3(\text{cnccNV}))))))
 \end{array}$$

5.5.4 Traces

Chunking builds a trace for each instantiation starting with the set of instantiated conditions of the instantiation's matching. The trace uses the contents of temporary memory to segregate the conditions into three classes: grounds, potentials and locals. It also records the not doubletons that the match constrained to be different and the instantiation's matches matching.

The grounds of a trace are the instantiated conditions that were directly matched from a parent goal through a chain of conditions. The potentials are those conditions that are in the transitive closure of a parent of the match goal, but were not directly matched from the parent. All of the non-ground negative conditions are also added into the potentials, as they can not be traced through because they never had a corresponding element in working memory. The locals are the remaining conditions, and should be exactly those positive conditions that are available only from the transitive closure of the match goal.

$Grounds : \mathbb{P} TME \times Instantiation \rightarrow \mathbb{P} Condition$
 $Potentials : \mathbb{P} TME \times Instantiation \rightarrow \mathbb{P} Condition$
 $Locals : \mathbb{P} TME \times Instantiation \rightarrow \mathbb{P}(\text{ran } PositiveCondition)$
 $Nots : Instantiation \rightarrow \mathbb{P} Not$

$\forall TM : \mathbb{P} TME; i : Instantiation \bullet$

$Grounds(TM, i) =$
 $TCICondition($
 $OfldCondition($
 $GoalAncestors(MatchGoal(i.match, TM), TM),$
 $\text{ran } i.matching), \text{ran } i.matching)$

$\forall TM : \mathbb{P} TME; i : Instantiation \bullet$

$Potentials(TM, i) =$
 $((\text{ran } i.matching \setminus \text{ran } PositiveCondition) \cup$
 $TCICondition($
 $OfldCondition($
 $(TCITME(TM))(GoalAncestors(MatchGoal(i.match, TM), TM)),$
 $\text{ran } i.matching), \text{ran } i.matching) \setminus Grounds(TM, i))$

$\forall TM : \mathbb{P} TME; i : Instantiation \bullet$

$Locals(TM, i) = \text{ran } i.matching \setminus (Grounds(TM, i) \cup Potentials(TM, i))$

$\forall i : Instantiation \bullet$

$Nots(i) = \bigcup \{c : \text{dom } i.matching \bullet ConditionsNots(i.binding.c, i.matching(c))\}$

The *Trace* data schema holds the grounds, potentials, locals, nots and a matching. The matching is a correspondence between the instantiated versions of the positive conditions and the working memory element instances that they matched.

$$\begin{array}{l}
 \text{Trace} \\
 \hline
 \text{grounds, potentials, locals} : \mathbb{P} \text{ Condition} \\
 \text{nots} : \mathbb{P} \text{ Not} \\
 \text{matching} : \text{Condition} \rightarrow (\text{WME} \times \mathbb{N}) \\
 \hline
 \text{dom matching} = (\text{grounds} \cup \text{potentials} \cup \text{locals}) \cap \text{ran PositiveCondition}
 \end{array}$$

The function *makeTrace* is provided to simplify the creation of traces by chunking.

$$\begin{array}{l}
 \text{makeTrace} : \mathbb{P} \text{ TME} \times \text{Instantiation} \rightarrow \text{Trace} \\
 \hline
 \forall TM : \mathbb{P} \text{ TME}; i : \text{Instantiation} \bullet \\
 \text{makeTrace}(TM, i) = \\
 \quad (\mu \text{ grounds} : \{\text{Grounds}(TM, i)\}; \\
 \quad \text{potentials} : \{\text{Potentials}(TM, i)\}; \\
 \quad \text{locals} : \{\text{Locals}(TM, i)\}; \\
 \quad \text{nots} : \{\text{Nots}(i)\}; \\
 \quad \text{matching} : \{i.\text{match}.\text{matching}\} \bullet \\
 \quad \theta \text{Trace})
 \end{array}$$

The components of a trace is the components of its conditions, nots and the components of the working memory element instances that it matched.

$$\begin{array}{l}
 \text{TracesComponents} : \text{Trace} \rightarrow \mathbb{P} \Sigma \\
 \hline
 \forall t : \text{Trace} \bullet \\
 \text{TracesComponents}(t) = \\
 \quad \cup (\text{ConditionsComponents}(t.\text{grounds} \cup t.\text{potentials} \cup t.\text{locals})) \cup \\
 \quad \cup t.\text{nots} \cup \\
 \quad \cup \text{first}(\text{ran matching})
 \end{array}$$

5.5.5 Chunking's State

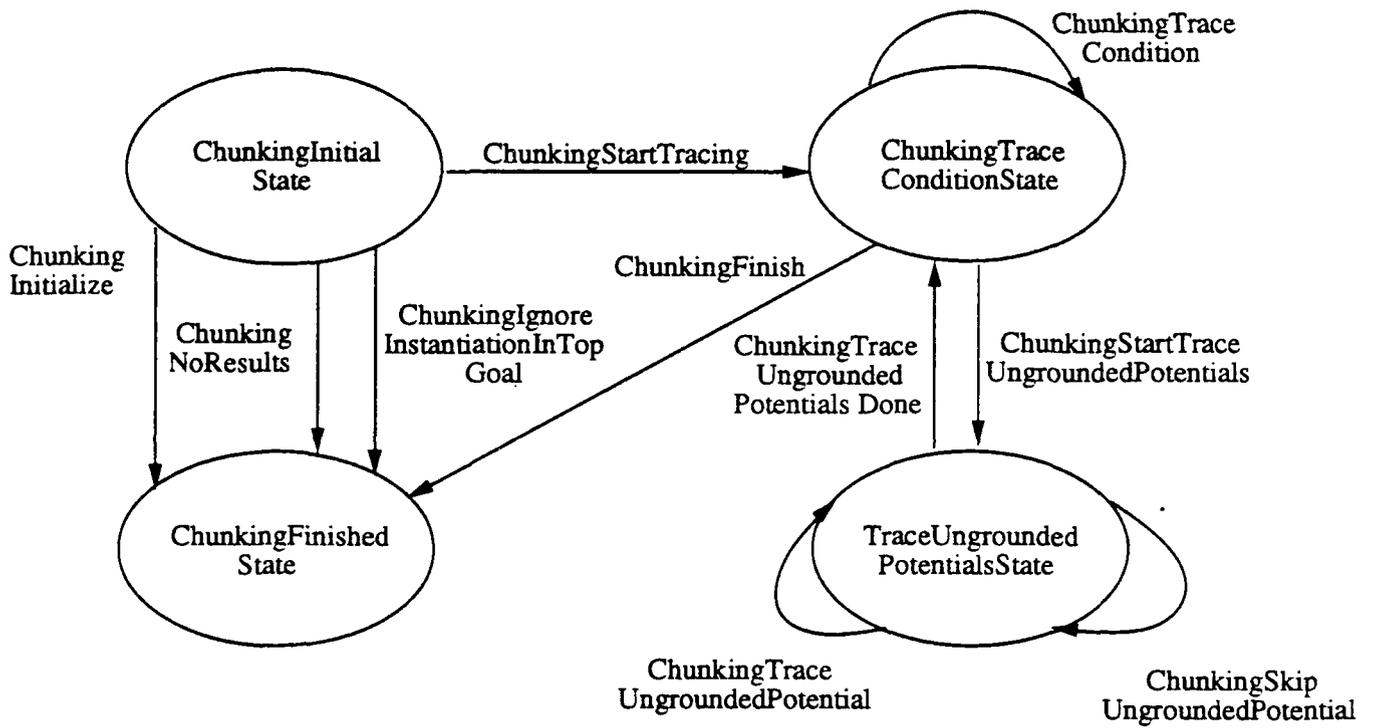


Figure 5.1: Chunking's State Machine

Chunking's state machine has only four states, but has the most complicated state schema of any of Soar's state machines.

- *ChunkingInitialState* — chunking starts in this state.
- *ChunkingTraceConditionState* — traces through the local conditions.
- *ChunkingTraceUnGroundedPotentialsState* — traces through all of the ungrounded potentials, one at a time.
- *ChunkingFinishedState* — chunking finishes in this state.

```
ChunkingState ::= ChunkingInitialState
                | ChunkingTraceConditionState
                | ChunkingTraceUnGroundedPotentialsState
                | ChunkingFinishedState
```

When chunking processes an instantiation, it may or may not produce a chunk, depending upon the results of the instantiation's right hand side. If it does produce a new chunk, then it returns the new instantiation for that chunk so that preference phase can process it. The *ChunkingResult* type is used to inform the preference phase that chunking is returning a new instantiation.

```
ChunkingResult ::= NewInstantiation NoInstantiation
```

Soar is actually viewed as continuously learning from its operation. However, in practice Soar's learning mechanism is the most difficult part of Soar to use. Chunking is difficult to use for a variety of reasons: it is rather complicated so users find it difficult to understand, it can create expensive rules which prohibitively degrade performance ([Tam91]), it has some known over-generalization problems ([LCAS90], 70, 73, 74), and it has yet to be well integrated with our theoretical framework — the problem space computational model ([NYL+91]).

For convenience we allow Soar users to turn learning partially off in two ways. The chunking state schema has a switch *learn* of type *Learn*. When *learn* is off or if chunking backtraces through a local goal's 'QUIESCENCE' element, a chunk is learned but it is not variabilized and not-ified. These internal chunks act like restricted productions; they can match only one sequence of elements, they confer support on their resulting preferences, and they are guaranteed to be very cheap to match. When an internal chunk's single instantiation retracts, Soar deletes the production from production memory as it is never likely to match again.

```
Learn ::= On | Off
```

Chunking shares six memories with other modules:

- *SPM* — the set of productions
- *TM* — temporary memory
- *WM* — working memory
- *PM* — preference memory
- *GM* — goal memory and
- *IS* — instantiation memory, the set of matches that have been instantiated by the preference phase.

Chunking

SPM : $\mathbb{P} SP$
TM : $\mathbb{P} TME$
WM : $\mathbb{P} WME$
PM : $\mathbb{P} Preference$
GM : $\mathbb{P} Identifier$
IS : $\mathbb{P} Instantiation$
WM# : $WME \mapsto \mathbb{N}$
PM# : $Preference \mapsto \mathbb{N}$
maz_preference, maz_wme : \mathbb{N}
TrP : $(Preference \times \mathbb{N}) \times Identifier \mapsto Trace$
TrW : $(WME \times \mathbb{N}) \times Identifier \mapsto (Preference \times \mathbb{N})$
learn : *Learn*
instantiation : *Instantiation*
goal : *Identifier*
results : $\mathbb{P} Preference$
seeds : $\mathbb{P} Preference$
grounds, potentials, ungrounded_potentials : $\mathbb{P} Condition$
locals : $\mathbb{P}(\text{ran } PositiveCondition)$
nots : $\mathbb{P} Not$
matching : $Condition \mapsto (WME \times \mathbb{N})$
chunking_state : *ChunkingState*
chunking_result : *ChunkingResult*

goal $\in GM$

$\text{dom } WM\# = WM$

$\text{dom } PM\# = PM$

$(first ; second)(TrP) \subseteq GM$

$(first ; second)(TrW) \subseteq GM$

Chunking introduces four new memories of indefinite extent.

- *TrP* — the trace memory for preferences.
- *TrW* — the trace memory for working memory elements.
- *WM #* — a numbering function that distinguishes the instances of each working memory element in memory.
- *PM #* — a numbering function that distinguishes the instances of each preference in memory.

WM # and *PM #* are partial injections that provide a way for chunking to distinguish between two different elements that were in memory at different times. The numberings were defined in this section because only chunking must actually distinguish between the instances of memory elements. As preference phase generates new preferences, and decide and IO generate new working memory elements they extend the numbering to give each new element a new number. We chose this numbering function representation instead of giving each working memory element and preference a number component because it allowed us to still use the standard set theoretic operations, instead of having to define our own little set theory for just sets of *WMEs* and *Preferences*.

TrP records for each preference, in each goal, the trace that generated that preference instance. Similarly, *TrW* records for each working memory element instance, in each goal, the acceptable or require preference instance that was in memory when decide added the working memory element instance. The invariants on the chunking schema require that the traces are only for current goals. The specification would be more accurate if the trace memories were explicitly contracted when goals are popped, but this simpler data invariant describes the intent without making the impasser's already large operations larger.

Chunking holds a dozen other pieces of state for use during its state machine's execution.

- instantiation — the instantiation that preference phase has passed to chunking.
- goal — the match goal of the instantiation.
- results — the results of the instantiation.
- seeds — the seeds, starts out as the set of results but is deleted as each result backtraced.
- grounds, potentials, locals — the grounds, potentials and locals set of chunking's search through the trace memories.
- nots — the nots of chunking's search.

- *matching* — the matching that maps chunking's conditions (grounds, potentials and locals) to the working memory instances they matched.
- *chunking_state* — the state counter.
- *chunking_result* — a flag that informs the preference phase if it is receiving a new instantiation back that it should calculate support from.
- *ungrounded_potentials* — the set of ungrounded potentials that chunking must backtrack one step.

5.5.6 Changing, Initializing and Resetting Chunking

When chunking runs, it does not modify the contents of temporary memory. If it changed temporary memory then the calculations for the traces of instantiation from the same parallel preference phase would have ordering dependencies. We define Δ *Chunking* to emphasize that chunking does not modify temporary memory and that it adds to the contents of production memory.

Δ <i>Chunking</i> <i>Chunking</i> <i>Chunking'</i>
$TM' = TM$
$SPM \subseteq SPM'$

InitChunking specifies legal initial states for chunking.

<i>InitChunking</i>
<i>Chunking</i>
$TrP = \emptyset$
$TrW = \emptyset$
$results = \emptyset$
$seeds = \emptyset$
$grounds = potentials = \emptyset$
$locals = \emptyset$
$nots = \emptyset$
$matching = \emptyset$
$chunking_state = ChunkingFinishedState$
$maz_preference = 0$
$maz_wme = 0$

ChunkingInitialize initializes chunking by placing it in the initial state.

<i>ChunkingInitialize</i>
Δ <i>Chunking</i>
$chunking_state = ChunkingFinishedState$
$results = \emptyset$
$seeds = \emptyset$
$grounds = potentials = \emptyset$
$locals = \emptyset$
$nots = \emptyset$
$matching = \emptyset$
$chunking_state' = ChunkingInitialState$

When chunking is externally reset, it must have all of its memories for the current search through trace memory emptied.

<i>ChunkingReset</i> Δ <i>Chunking</i>
<i>chunking_state'</i> = <i>ChunkingFinishedState</i>
<i>results</i> = \emptyset
<i>seeds</i> = \emptyset
<i>grounds</i> = <i>potentials</i> = \emptyset
<i>locals</i> = \emptyset
<i>nots</i> = \emptyset
<i>matching</i> = \emptyset

5.5.7 Defining Results

When an instantiation modifies the transitive closure of a goal, it is said to add results to that goal. The results are the preferences that the instantiation would make available when its right hand side's preferences are added in isolation to preference memory. These result preferences are used as the starting point, or seeds, of chunking's search through trace memory.

$$\begin{aligned}
 & \text{Results} : \text{Identifier} \times \text{Instantiation} \times \mathcal{P} \text{ TME} \rightarrow \mathcal{P} \text{ Preference} \\
 & \forall g : \text{Identifier}; i : \text{Instantiation}; TM : \mathcal{P} \text{ TME} \bullet \\
 & \quad \exists TMplusRHS : \mathcal{P} \text{ TME} \quad TMplusRHS = TM \cup \text{Preference TME} \{i.rhs\} \bullet \\
 & \quad \text{Results}(g, i, TM) = \\
 & \quad \quad \text{Preference TME} \sim \{(\text{OfId}_{\text{TME}}(\text{TCI}_{\text{TME}}(TMplusRHS))(\{g\}), TMplusRHS) \\
 & \quad \quad \quad \setminus (\text{OfId}_{\text{TME}}(\text{TCI}_{\text{TME}}(TM))(\{g\}), TM)\} \cap \text{ran Preference TME}
 \end{aligned}$$

5.5.8 Chunking an Instantiation

Chunking does one of three things with each instantiation: starts tracing it, skips it because it is in the top goal, or skips it because it generates no results.

If the instantiation has results, *ChunkingStartTracing* makes a trace for it, and initializes the search through the backtrace memories from the trace's components.

ChunkingStartTracing

Δ *Chunking*

chunking_state = *ChunkingInitialState*
goal' = *MatchGoal*(*instantiation.match*, *TM*)
TrP' = *TrP* \oplus
 { *p* : *instantiation.rhs* •
 ((*p.PM* $\#$ (*p*)), *goal*) \rightarrow *makeTrace*(*TM*, *instantiation*)}
 $\exists g$: *GM* •
makeWME(*goal'*, 'OBJECT', *g*, *No*) \in *WM* \wedge
Results(*g*, *instantiation.TM*) $\neq \emptyset$ •
(*results'* = *Results*(*g*, *instantiation.TM*) \wedge
seeds' = *results'*)
grounds' = *Grounds*(*TM*, *instantiation*)
potentials' = *Potentials*(*TM*, *instantiation*)
locals' = *Locals*(*TM*, *instantiation*)
nots' = *Nots*(*instantiation*)
matching' = *instantiation.match.matching*
chunking_state' = *ChunkingTraceConditionState*

If the match goal of the chunk is the top goal, then it can not give any results to a parent goal. This makes it uninteresting to the chunker, so it does not need to store a trace. Chunking finishes and sets the results flag to tell preference phase that no new instantiation was generated.

ChunkingIgnoreInstantiationInTopGoal

Δ *Chunking*

chunking_state = *ChunkingInitialState*
 $\neg (\exists g$: *GM* •
makeWME(*MatchGoal*(*instantiation.match*, *TM*), 'OBJECT', *g*, *No*)
 \in *WM*)
chunking_result' = *NoInstantiation*
chunking_state' = *ChunkingFinishedState*

If the trace has no results, then chunking builds the trace for its preferences and returns to preference phase.

<p><i>ChunkingNoResults</i></p> <p>Δ <i>Chunking</i></p> <p><i>chunking_state</i> = <i>ChunkingInitialState</i></p> <p>$\exists g : GM \mid \text{makeWME}(\text{goal}, \text{'OBJECT'}, g, No) \in WM \bullet$ $\text{Results}(g, \text{instantiation}, TM) = \emptyset$</p> <p>$\text{Tr}P' = \text{Tr}P \oplus$ $\{p : \text{instantiation.rhs} \bullet$ $((p, PM\#(p)), \text{goal}) \mapsto \text{makeTrace}(TM, \text{instantiation})\}$</p> <p><i>chunking_result'</i> = <i>NoInstantiation</i></p> <p><i>chunking_state'</i> = <i>ChunkingFinishedState</i></p>

When Decide fills out slots in working memory, it must augment the *TrW*. It gives each element a trace that maps it, under the oldest goal that it is in the transitive closure of, to the require or acceptable preference instance for it in preference memory. The *Owner* function finds the oldest goal that an identifier is in.

<p><i>Owner</i> : <i>Identifier</i> \times \mathbb{P} <i>Identifier</i> \times \mathbb{P} <i>TME</i> \mapsto <i>Identifier</i></p> <p>$\forall i : \text{Identifier}; GM : \mathbb{P} \text{Identifier}; TM : \mathbb{P} \text{TME} \bullet$ $\exists g : GM \mid i \in \text{TCM}(g, TM) \wedge$ $\neg (\exists \text{anc} : \text{GoalAncestors}(g, TM) \bullet i \in \text{TCM}(\text{anc}, TM) \bullet$ $\text{Owner}(i, GM, TM) = g$</p>

TraceWME provides an operation for decide and io to use to augment the working memory instance number and trace. As preference semantics processes requires before acceptables, *TraceWME* favors requires over acceptable preferences.

<p><i>TraceWME</i></p> <p>Δ <i>Chunking</i></p> <p>$w : WME$</p> <hr/> <p>$maz_wme' = maz_wme + 1$</p> <p>$WM\# = WM\# \oplus \{w \mapsto maz_wme'\}$</p> <p>$\exists p : PM \mid$ $p = makeUnaryPreference(first(w?).id, first(w?).attribute, first(w?).value, '!') \bullet$ $TrW' = Trw \oplus \{(w, WM\#(w)), Owner(w.id) \mapsto (p, PM\#(p))\}$</p> <p>$\neg (\exists p : PM \bullet$ $p = makeUnaryPreference(w.id, w.attribute, w.value, '!'))$</p> <p>$\wedge$</p> <p>$\exists p : PM \mid$ $p = makeUnaryPreference(w.id, w.attribute, w.value, '+') \bullet$ $TrW' = Trw \oplus \{(w?, Owner(w.id)) \mapsto (p, PM\#(p))\}$</p>

TraceItems adds new traces for the items augmentation of impasses. The items augmentations are traced back to the require or acceptable preference that caused them to be considered for their slot.

<p><i>TraceItems</i></p> <p>Δ <i>Chunking</i></p> <p>$W : P WME$</p> <p>$i : Identifier$</p> <hr/> <p>$maz_wme' = maz_wme - \#W$</p> <p>$\exists f : W \mapsto (maz_wme + 1 .. maz_wme') \bullet WM\# = WM\# \oplus f$</p> <p>$\exists object, attribute : Identifier \mid$ $makeWME(i, 'OBJECT', object) \in WM \wedge$ $makeWME(i, 'ATTRIBUTE', attribute) \in WM \bullet$ $TrW' = Trw \oplus$ $\{w : W; p : PM \mid p = makeUnaryPreference(object, attribute, w.value, '+') \bullet$ $((w, WM\#(w)), Owner(w.id)) \mapsto (p, PM\#(p))\} \oplus$ $\{w : W; p : PM \mid p = makeUnaryPreference(object, attribute, w.value, '!') \bullet$ $((w, WM\#(w)), Owner(w.id)) \mapsto (p, PM\#(p))\}$</p>
--

5.5.9 Tracing Seeds, Locals and Grounded Potentials

The *ChunkingTraceConditionState* performs most of the search through trace memory. Its one composite operation, *ChunkingTraceCondition*, traces through local conditions in one of seven ways.

- *ChunkingTraceSeed* — if a seed has a preference trace for the goal, then follow it.
- *ChunkingSkipSeed* — if a seed does not have a preference trace for the goal, then skip it.
- *ChunkingTraceLocal* — if a local's working memory element has a trace, then follow it.
- *ChunkingSkipArchitectureGeneratedLocal* — if a local matched an architecturally generated working memory element, skip it.
- *ChunkingTraceQuiescenceLocal* — if a local matched the goal's quiescence working memory element then add a quiescence condition to the grounds set.
- *ChunkingPotentializeLocal* — if a local condition's working memory element does not have a working memory element trace for this goal, then move it to the potential's set.
- *ChunkingTraceGroundedPotential* — if there are no more locals, and a potential's condition is in the transitive closure of the grounds set, then move it to the grounds.

The seed preferences are the starting point for chunking's search through trace memory. The seeds are processed first, and then the locals and finally the potentials. If a seed has a preference memory trace, chunking follows it by removing the seed from the seed set, and unioning the trace's components into it's search memories.

<p><i>ChunkingTraceSeed</i></p> <hr/> <p>Δ <i>Chunking</i></p> <hr/> <p>$chunking_state = ChunkingTraceConditionState := chunking_state'$</p> <p>$\exists s : seeds \bullet$</p> <p style="padding-left: 20px;">$\exists t : Trace \mid t = TrP((s, PM\#(s)), goal) \bullet$</p> <p style="padding-left: 40px;">$((seeds' = seeds \setminus \{s\}) \wedge$</p> <p style="padding-left: 40px;">$(grounds' = grounds \cup t.grounds) \wedge$</p> <p style="padding-left: 40px;">$(potentials' = potentials \cup t.potentials) \wedge$</p> <p style="padding-left: 40px;">$(locals' = locals \cup t.locals) \wedge$</p> <p style="padding-left: 40px;">$(nots' = nots \cup t.not) \wedge$</p> <p style="padding-left: 40px;">$(matching' = matching \oplus t.matching))$</p>
--

Chunking skips seeds that do not have a preference trace in the match goal. They were generated in the parent goal and are in the transitive closure of the results without being part of the results.

$\begin{array}{l} \text{ChunkingSkipSeed} \\ \Delta \text{Chunking} \\ \text{chunking_state} = \text{ChunkingTraceConditionState} = \text{chunking_state}' \\ \exists s : \text{seeds} \mid \\ \neg (\exists t : \text{Trace} \mid t = \text{TrP}((s, \text{PM} \#(s)), \text{goal})) \bullet \\ ((\text{seeds}' = \text{seeds} \setminus \{s\}) \wedge \\ (\text{grounds}' = \text{grounds}) \wedge \\ (\text{potentials}' = \text{potentials}) \wedge \\ (\text{locals}' = \text{locals}) \wedge \\ (\text{nots}' = \text{nots}) \wedge \\ (\text{matching}' = \text{matching})) \end{array}$

Chunking's basic step in the search is to backtrack through a local. The local condition's matched working memory element is chased through the working memory trace, and the resulting preference instance is chased through the preference trace to arrive at a trace. This trace's components are unioned into the search memories, and the local is discarded.

$\begin{array}{l} \text{ChunkingTraceLocal} \\ \Delta \text{Chunking} \\ \text{chunking_state} = \text{ChunkingTraceConditionState} = \text{chunking_state}' \\ \text{seeds} = \emptyset \\ \exists l : \text{locals} \bullet \\ \exists t : \text{Trace} \mid \\ t = \text{TrP}(\text{TrW}(\text{matching}(l), \text{goal}), \text{goal}) \bullet \\ ((\text{grounds}' = \text{grounds} \cup t.\text{grounds}) \wedge \\ (\text{potentials}' = \text{potentials} \cup t.\text{potentials}) \wedge \\ (\text{locals}' = (\text{locals} \cup t.\text{locals}) \setminus \{l\}) \wedge \\ (\text{nots}' = \text{nots} \cup t.\text{nots}) \wedge \\ (\text{matching}' = (\{l\} \triangleleft \text{matching}) \oplus t.\text{matching})) \end{array}$

When decide constructs a goal or an impasse it represents it as elements in working memory. *ImpasseAttribute* is the set of the attributes it uses to represent goals and impasses.

ImpasseAttribute : P Constant

ImpasseAttribute =
 {`PROBLEM-SPACE`, `STATE`, `OPERATOR`, `OBJECT`,
 `ATTRIBUTE`, `IMPASSE`, `CHOICES`, `ITEM`, `QUIESCENCE`,
 `TYPE`}

All but the items working memory elements of these goals have no working memory traces. Decide adds working memory traces that make the items depend upon the preference instance for the require or acceptable preference that caused decide to generate the impasse.

ChunkingSkipArchitectureGeneratedLocal _____

Δ *Chunking*

chunking_state = *ChunkingTraceConditionState* = *chunking_state'*

seeds = \emptyset

$\exists l : \text{locals}; w : \text{WME} :$

$w = \text{first}(\text{matching}(l)) \wedge$

$w.\text{id} \in \text{GM} \wedge$

$w.\text{attribute} \in \text{ImpasseAttribute} \setminus \{\text{ITEM}\} \wedge$

$w.\text{context_acceptable_preference} = \text{No} \bullet$

$((\text{grounds}' = \text{grounds}) \wedge$

$(\text{potentials}' = \text{potentials}) \wedge$

$(\text{locals}' = \text{locals} \setminus \{l\}) \wedge$

$(\text{nots}' = \text{nots}) \wedge$

$(\text{matching}' = \{l\} \triangleleft \text{matching}))$

The 'QUIESCENCE' goal working memory element is used to turn off chunking for a goal that the user does not want to chunk. When the user knows that the processing in a subgoal will generate a problem, she simply makes one or more of the productions that fire in the subgoal match the goal's quiescence working memory element. When chunking backtraces through this quiescence element, *ChunkingTraceQuiescenceLocal* traces it back to the quiescence element of the parent goal. When chunking has finished, the presence of the quiescence augmentation in the grounds set tells chunking not to variabilize the new chunk.

ChunkingTraceQuiescenceLocal

Δ *Chunking*

chunking_state = *ChunkingTraceConditionState* = *chunking_state'*

seeds = \emptyset

$\exists l : \text{locals}; w : \text{WME}; g : \text{GM} \mid$

$w = \text{first}(\text{matching}(l)) \wedge$

$w.\text{id} \in \text{GM} \wedge$

$w.\text{attribute} = \text{'QUIESCENCE'} \wedge$

$w.\text{value} = \text{'T'} \wedge$

$w.\text{context_acceptable_preference} = \text{No} \wedge$

$\text{makeWME}(\text{goal}, \text{'OBJECT'}, g, \text{No}) \in \text{WM} \bullet$

$\exists qc :$

$\{ \text{PositiveCondition}(\$

$\text{makeWMETest}(\text{EqualityTest}(g),$

$\text{EqualityTest}(\text{'QUIESCENCE'}),$

$\text{EqualityTest}(\text{'T'}, \text{NoTest})) \bullet$

$\exists qw : \text{WME} \mid$

$qw.\text{id} = g \wedge qw.\text{attribute} = \text{'QUIESCENCE'} \wedge$

$qw.\text{value} = \text{'T'} \wedge qw.\text{context_acceptable_preference} = \text{No} \bullet$

$((\text{grounds}' = \text{grounds} \cup \{qc\}) \wedge$

$(\text{potentials}' = \text{potentials}) \wedge$

$(\text{locals}' = \text{locals} \setminus \{l\}) \wedge$

$(\text{nots}' = \text{nots}) \wedge$

$(\text{matching}' = (\{l\} \Leftarrow \text{matching}) \oplus \{qc \mapsto (qw, \text{WM}\#(qw))\}))$

When chunking traces back to a local that does not have an trace record for the current goal, chunking makes the local a potential. A working memory element that was accessible locally in a match but does not have trace for that goal, has to have been brought down from a parent goal. For this local to be accessible in the current goal, but not to have a trace record for the goal, it must have been brought down as part of the transitive closure of an element from a parent goal, which later had its parent goal link snapped. So chunking adds it to the set of potentials as there will be a chain of conditions through trace memory that will eventually ground out this local.

ChunkingPotentializeLocal

Δ *Chunking*

$chunking_state = ChunkingTraceConditionState = chunking_state'$
 $seeds = \emptyset$
 $\exists l : locals \mid$
 $\neg (\exists t : Trace \bullet t = TrP(TrW(matching(l), goal), goal)) \bullet$
 $((grounds' = grounds) \wedge$
 $(potentials' = potentials \cup \{l\}) \wedge$
 $(locals' = locals \setminus \{l\}) \wedge$
 $(nots' = nots) \wedge$
 $(matching' = \{l\} \Leftarrow matching))$

When there are no more seeds and locals to process, chunking adds the potentials that are connected to the grounds into the grounds set.

ChunkingTraceGroundedPotential

Δ *Chunking*

$chunking_state = ChunkingTraceConditionState = chunking_state'$
 $seeds = \emptyset$
 $locals = \emptyset$
 $\exists gp : potentials \mid$
 $gp \in TCI_{Condition}(grounds, grounds \cup potentials) \cap potentials \bullet$
 $((grounds' = grounds \cup \{gp\}) \wedge$
 $(potentials' = potentials \setminus \{gp\}) \wedge$
 $(locals' = \emptyset) \wedge$
 $(nots' = nots) \wedge$
 $(matching' = \{gp\} \Leftarrow matching))$

ChunkingTraceCondition is a composite operation that groups the tracing for locals and grounded potentials.

$$\begin{aligned}
 \textit{ChunkingTraceCondition} &\equiv \\
 &\textit{ChunkingTraceSeed} \vee \textit{ChunkingSkipSeed} \vee \\
 &\textit{ChunkingTraceLocal} \vee \textit{ChunkingSkipArchitectureGeneratedLocal} \vee \\
 &\textit{ChunkingTraceQuiescenceLocal} \vee \textit{ChunkingPotentializeLocal} \vee \\
 &\textit{ChunkingTraceGroundedPotential}
 \end{aligned}$$

5.5.10 Tracing UnGrounded Potentials

Chunking is not able to connect all of the potential conditions into the grounds set. When the seeds and locals are empty, and no potentials that are connected to the grounds set remains, *ChunkingStartTraceUnGroundedPotentials* moves chunking into the trace ungrounded potentials state. This state backtraces one step through all of the positive potentials. It saves all of the entering potentials in the *ungrounded_potentials* set, so that it can augment the potentials set with the potentials of newly backtraced traces.

$$\begin{array}{l}
 \textit{ChunkingStartTraceUnGroundedPotentials} \\
 \Delta \textit{Chunking} \\
 \hline
 \textit{chunking_state} = \textit{ChunkingTraceConditionState} \\
 \textit{seeds} = \emptyset \\
 \textit{locals} = \emptyset \\
 - (\exists gp : \textit{potentials} \bullet \\
 \quad gp \in \textit{TCICondition}(\textit{grounds}, \textit{grounds} \cup \textit{potentials}) \cap \textit{potentials}) \\
 \textit{potentials}' = \emptyset \\
 \textit{ungrounded_potentials}' = \textit{potentials} \\
 \textit{chunking_state}' = \textit{ChunkingTraceUnGroundedPotentialsState}
 \end{array}$$

Each positive potential with a trace is backtraced in the standard way.

$$\begin{array}{l}
 \text{--- } \textit{ChunkingTraceUnGroundedPotential} \text{ ---} \\
 \Delta \textit{Chunking} \\
 \text{---} \\
 \textit{chunking_state} = \textit{ChunkingTraceUnGroundedPotentialsState} = \textit{chunking_state}' \\
 \exists gp : \textit{ungrounded_potentials} \cap (\textit{ran PositiveCondition}) \bullet \\
 \quad \exists t : \textit{Trace} \mid t = \textit{TrP}(\textit{TrW}(\textit{matching}(gp), \textit{goal}), \textit{goal}) \bullet \\
 \quad ((\textit{grounds}' = \textit{grounds} \cup t.\textit{grounds}) \wedge \\
 \quad (\textit{ungrounded_potentials}' = \textit{ungrounded_potentials} \setminus \{gp\}) \\
 \quad (\textit{potentials}' = \cup t.\textit{potentials}) \wedge \\
 \quad (\textit{locals}' = \textit{locals} \cup t.\textit{locals}) \wedge \\
 \quad (\textit{nots}' = \textit{nots} \cup t.\textit{nots}) \wedge \\
 \quad (\textit{matching}' = \{gp\} \triangleleft \textit{matching})) \\
 \text{---}
 \end{array}$$

Chunking skips positive potentials that don't have a trace.

$$\begin{array}{l}
 \text{--- } \textit{ChunkingSkipUnGroundedPotential} \text{ ---} \\
 \Delta \textit{Chunking} \\
 \text{---} \\
 \textit{chunking_state} = \textit{ChunkingTraceUnGroundedPotentialsState} = \textit{chunking_state}' \\
 \exists gp : \textit{ungrounded_potentials} \cap (\textit{ran PositiveCondition}) : \\
 \quad \neg (\exists t : \textit{Trace} \bullet t = \textit{TrP}(\textit{TrW}(\textit{matching}(gp), \textit{goal}), \textit{goal})) \bullet \\
 \quad ((\textit{grounds}' = \textit{grounds}) \wedge \\
 \quad (\textit{ungrounded_potentials}' = \textit{ungrounded_potentials} \setminus \{gp\}) \\
 \quad (\textit{potentials}' = \textit{potentials}) \wedge \\
 \quad (\textit{locals}' = \textit{locals}) \wedge \\
 \quad (\textit{nots}' = \textit{nots}) \wedge \\
 \quad (\textit{matching}' = \textit{matching})) \\
 \text{---}
 \end{array}$$

When all of the positive potentials have been traced, chunking returns to the trace condition state to backtrace the newly generated locals. The potentials collected while backtracing the positive ungrounded potentials are unioned with the remaining negated ungroundable potentials.

ChunkingTraceUnGroundedPotentialsDone _____

Δ *Chunking*

chunking_state = *ChunkingTraceUnGroundedPotentialsState*

ungrounded_potentials \

 { *p* : *ungrounded_potentials* :

$\neg (\exists t : \text{Trace} \bullet t = \text{TrP}(\text{TrW}(\text{matching}(p), \text{goal}), \text{goal}))$ } \cap

 (*ran PositiveCondition*) = \emptyset

potentials' = *potentials* \cup *ungrounded_potentials*

ungrounded_potentials' = \emptyset

chunking_state' = *ChunkingTraceConditionState*

5.5.11 ChunkSP

When the seeds, the locals and all of the positive potentials and the groundable negative potentials have been traced, backtracing is done and chunking generates the new production. If learning is on and the grounds do not match a quiescence test, then chunking variabilizes the grounds and results to builds the new production. Chunking also constructs the match that the matcher would generate, and uses this in constructing an instantiation for the new chunk, which it adds to the instantiation set. The new instantiation is handed back to the preference phase. Preference phase calculates the new instantiation's support, and then sends it back into chunking to have its results chunked up the hierarchy of goals.

ChunkingFinishLearnOn

Δ *Chunking*

chunking_state = *ChunkingTraceConditionState*

seeds = \emptyset

locals = \emptyset

potentials \cap *ran PositiveCondition* = \emptyset

learn = *On*

$\neg (\exists qc : \text{PositiveCondition} \sim (\text{grounds} \cap \text{ran PositiveCondition}) \bullet$
EqualityTest $\sim (qc.id) \in GM \wedge$
EqualityTest $\sim (qc.attribute) \wedge = \text{'QUIESCENCE'} \wedge$
EqualityTest $\sim (qc.value) = \text{'T'}$)

TCICondition(*grounds*, *grounds* \cup *potentials*) \cap *potentials* = \emptyset

$\exists a : \text{Assignment}; \text{name} : \text{Symbol} \bullet$
 $\exists \text{negativeconditions} : \{ \text{VariabilizeSetOfConditions}(\text{grounds} \setminus \text{ran PositiveCondition}, a) \} \bullet$
 $\exists \text{cwcc} : \{ \text{NotifyConditionMatchings}(\text{VariabilizeConditionMatching}(\text{grounds} \triangleleft \text{matching}, a), \text{NotsToVNots}(\text{nots}, a), \emptyset) \} \bullet$
 $\exists \text{match_matching} : \{ \text{first}(\text{first}_3(\text{cwcc})) \} \bullet$
 $\exists \text{instantiation_matching} : \{ \text{second}(\text{first}_3(\text{cwcc})) \} \bullet$
 $\exists sp : SP \mid sp = \text{makeSP}(\text{name}, \text{negativeconditions} \cup \text{dom match_matching}, \text{VariabilizeRHS}(\text{results}, a)) \bullet$
 $\exists m : \text{Match} \mid m = \text{makeMatch}(sp, \text{match_matching}) \bullet$
 $\exists i : \text{Instantiation} \mid i = \text{makeInstantiation}(m, \text{instantiation_matching}) \bullet$
 $(SPM' = SPM \cup \{sp\} \wedge$
 $IS' = IS \cup \{i\})$

chunking_result = *NewInstantiation*

chunking_state' = *ChunkingFinishedState*

If learning is off or quiescence is matched, chunking creates the new match and instantiation without variabilizing it and returns the instantiation to preference phase.

<p><i>ChunkingFinishLearnOff</i> _____</p> <p>Δ <i>Chunking</i></p> <hr/> <p><i>chunking_state</i> = <i>ChunkingTraceConditionState</i></p> <p><i>learn</i> = <i>Off</i> \vee</p> <p>($\exists qc : PositiveCondition \sim \{grounds \cap ran\ PositiveCondition\} \bullet$ <i>EqualityTest</i>$\sim(qc.id) \in GM \wedge$ <i>EqualityTest</i>$\sim(qc.attribute) \wedge = 'QUIESCENCE' \wedge$ <i>EqualityTest</i>$\sim(qc.value) = 'T'$)</p> <p>$\exists m : Match \mid m = makeMatch(sp, grounds \triangleleft matching) \bullet$ $\exists i : Instantiation \mid i = makeInstantiation(m, instantiation_matching) \bullet$ (<i>SPM'</i> = <i>SPM</i> $\cup \{sp\} \wedge$ <i>IS'</i> = <i>IS</i> $\cup \{i\}$)</p> <p><i>chunking_state'</i> = <i>ChunkingFinishedState</i></p>

5.5.12 Step Chunking

ChunkingStep groups all of the operations of chunking together into one operation that drives the entire state machine.

ChunkingStep $\hat{=}$

ChunkingStartTracing \vee *ChunkingIgnoreInstantiationInTopGoal* \vee
ChunkingNoResults \vee *ChunkingTraceCondition* \vee
ChunkingStartTraceUnGroundedPotentials \vee
ChunkingTraceUnGroundedPotential \vee *ChunkingSkipUnGroundedPotential* \vee
ChunkingTraceUnGroundedPotentialsDone \vee
ChunkingFinishLearnOn \vee *ChunkingFinishLearnOff*

5.6 Slots

The modules of Soar all view preference and working memory, as being partitioned into sets of preferences for individual slots. Each slot is an identifier, symbol pair.

Slot == *Identifier* \times *Symbol*

The preferences for a slot are those that hold the identifier of the slot in their id component, and the symbol of the slot in their attribute component.

Similarly, it views working memory as partitioned into sets of working memory elements for slots. The slot's working memory elements are those that hold the slot's identifier in their id component, and the symbol in their attribute, and are not context acceptable preferences.

$$\begin{array}{l}
 \vdots \text{ SlotsWMEs} : \text{Slot} \times \mathcal{P} \text{ WME} \rightarrow \mathcal{P} \text{ WME} \\
 \hline
 \forall \text{ slot} : \text{Slot}; \text{ WM} : \mathcal{P} \text{ WME} \bullet \\
 \text{SlotsWMEs}(\text{slot}, \text{WM}) = \\
 \{w : \text{WME} \mid w.\text{id} = \text{first}(\text{slot}) \wedge \\
 w.\text{attribute} = \text{second}(\text{slot}) \wedge \\
 w.\text{context_acceptable_preference} = \text{No}\}
 \end{array}$$

5.7 Recognition Memory's State

Recognition memory's state schema contains a state counter for preference phase and 12 memories, some of which are shared with other modules and some of which are local to recognition memory.

The preference phase has seven states.

1. *PPhaseInitialState* — the preference phase starts execution from here.
2. *PPhaseRetractInstantiationState* — this state removes the instantiations from the instantiation set which no longer have a match that is consistent with working memory.
3. *PPhaseInstantiateMatchState* — this state instantiates matches that do not yet have an instantiation.
4. *PPhaseConferSupportState* — after a new match has been instantiated, this state calculates the right hand side preference's support, adds the support to the support memories and collects the O-supported reject preferences.
5. *PPhaseChunkingState* — preference phase steps the execution of chunking from this state.
6. *PPhaseORejectState* — any new reject preferences with O-support are processed to remove matching preferences from preference memory and the support memories.
7. *PPhaseChangePMState* — preference phase changes preference memory to match support memory.
8. *PPhaseFinishedState* — when preference phase finishes, it rests in this state.

```
PreferencePhaseState ::= PPhaseInitialState  
| PPhaseRetractInstantiationState  
| PPhaseInstantiateMatchState  
| PPhaseConferSupportState  
| PPhaseChunkingState  
| PPhaseORejectState  
| PPhaseChangePMState  
| PPhaseFinishedState
```

Recognition Memory's state shares six memories with other modules of Soar.

- *WM* — the set of working memory elements. *WM* is shared with IO and Decide.
- *PM* — the set of preferences. *PM* is shared with Decide.
- *TM* — essentially a union of working memory and preference memory. *TM* is shared with Decide.

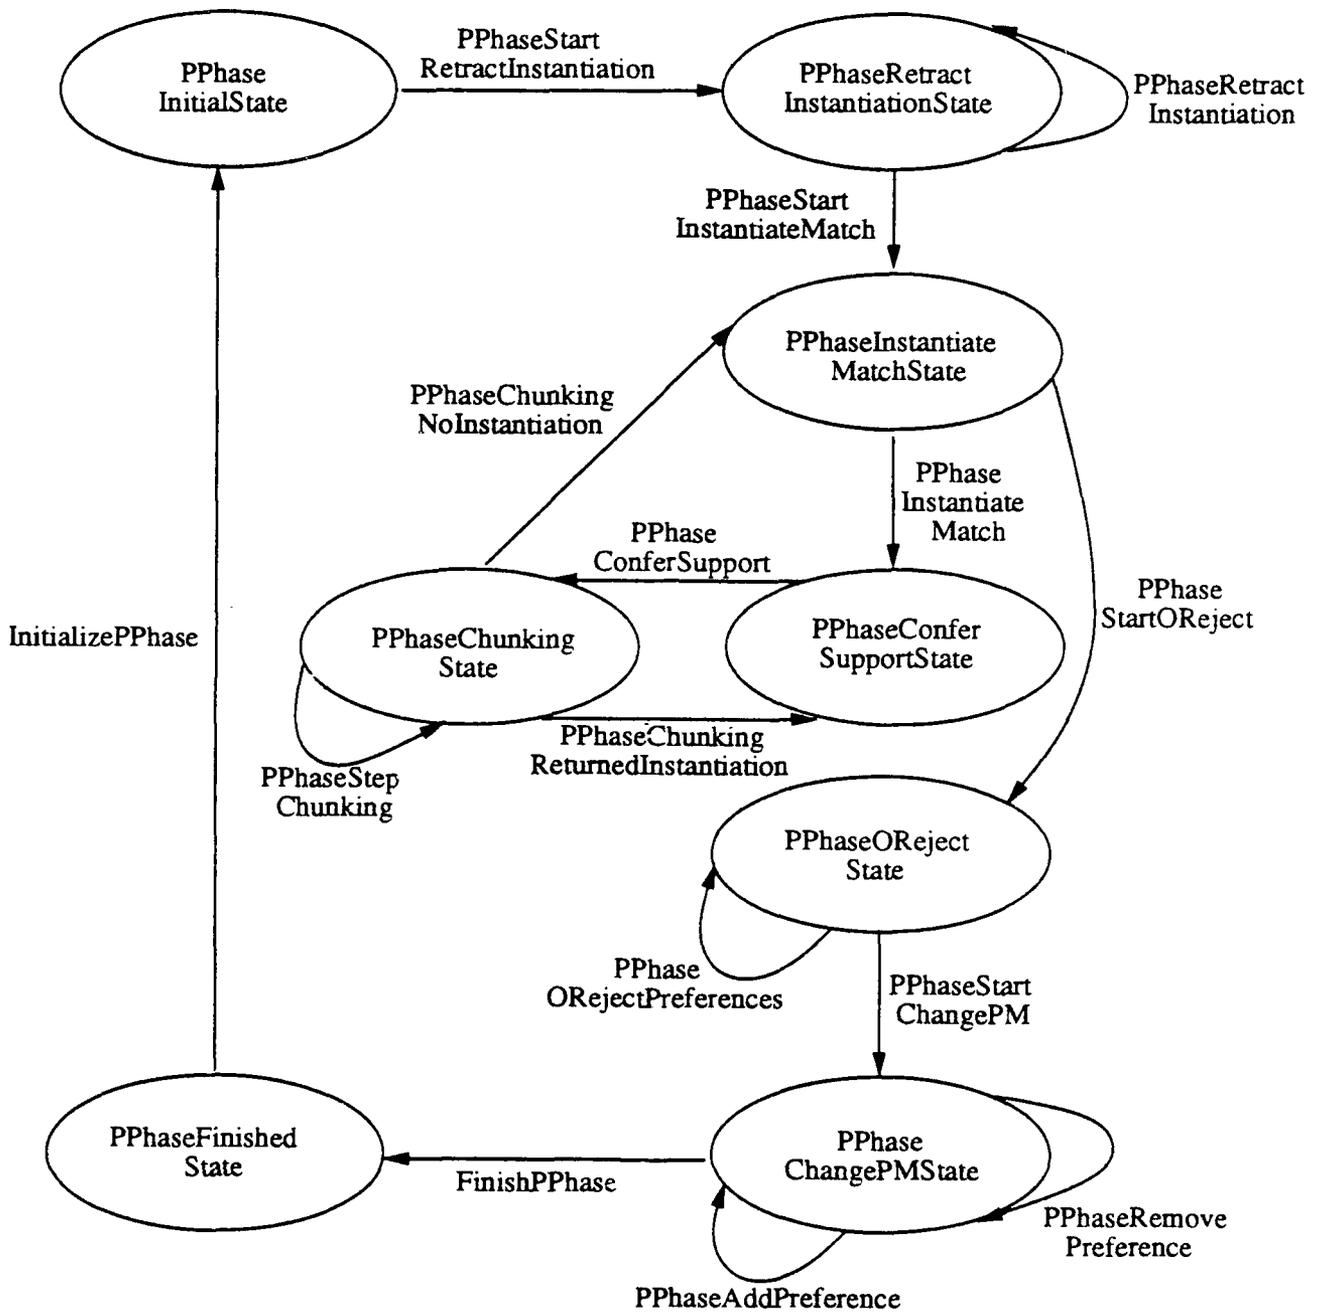


Figure 5.2: Preference Phase State Machine

- *IdentifierTable* — the set of identifiers in use in any part of Soar. The *IdentifierTable* is global to all modules of Soar.
- *GM* — is the set of goal identifiers. *GM* is read from Decide.
- *changed_slots* — the set of slots changed by preference phase. This is read and deleted by decide to tell which slots decide should consider.

RM

WM : P *WME*
PM : P *Preference*
TM : P *TME*
IdentifierTable : P *Identifier*
GM : P *Identifier*
SPM : P *SP*
MS : P *Match*
IS : P *Instantiation*
ISM : *Preference* \leftrightarrow *Instantiation*
OSM : *Preference* \rightarrow *Identifier*
ORM : P *Preference*
instantiation : *Instantiation*
pphase_state : *PreferencePhaseState*
changed_slots : P *Slot*
Chunking

$MS = \{m : Match \mid m.production \in SPM \wedge \text{Matches}_{WM}^{Match}(m, \langle VM, WM\# \rangle) \wedge m.binding \setminus \{m.production.roots\} \subseteq GM \cup IM\}$

$\text{ran } OSM \subseteq GM$

$SPM \neq \emptyset$

$TM = PreferenceTME \langle PM \rangle \cup WMETME \langle WM \rangle$

$\forall p : ORM \bullet$

$p \in \text{ran } UnaryP \wedge (UnaryP \sim (p)).preference = \text{'-'}'$

Recognition memory holds six local memories.

- *SPM* — is the set of productions in Soar.
- *MS* — the set of current matches for the productions. Each match must be for a production in production memory, must match the current contents of working memory, and all of the production's root variables must be bound to the identifiers of goals or impasses.
- *IS* — the set of instantiations that preference phase has created from the matches.

- *ISM* — the I-support memory. This memory is a relation from preferences to the instantiations that are currently supporting them.
- *OSM* — the O-support memory. This memory is a relation from preferences to the identifier of existing goals that had an instantiation I-support them.
- *ORM* — a memory of the O-supported reject preferences that preference phase has created while instantiating matches that have yet to be processed.

When recognition memory is initialized, all of its memories are empty except for the production memory. The production memory must not be empty so that Soar has some long term memories to use in performing its task; otherwise Soar would impasse indefinitely.

InitRM

RM

pphase_state = *PPhaseFinishedState*

IS = \emptyset

ISM = \emptyset

OSM = \emptyset

ORM = \emptyset

WM = \emptyset

PM = \emptyset

SPM = \emptyset

5.8 Support

Support is Soar's basic mechanism for determining the extent of preferences in preference memory. Support comes in two flavors: instantiation support, I-support, and operator support, O-support. Whenever an instantiation is created, it gives instantiation support to each of the preferences in its right hand side. When the instantiation is retracted, it withdraws its instantiation support from each of its preferences. Whenever instantiations that are creating, modifying or applying operators create a preference whose identifier is in the transitive closure of their match goal's states or operators, they confer operator support to these preference. The operator support has the extent of the match goal of the instantiation, and so is not retracted when the instantiation retracts.

This section defines the operators matched by an instantiation, when an instantiation is creating, modifying or applying an operator, and then combines these into a definition of when an instantiation O-supports a preference.

An instantiation matches an operator if it matches a goal's operator slot or operator acceptable preference.

$$\text{InstantiationsLHSOperators} : \text{Instantiation} \times \mathbb{P} \text{TME} \rightarrow \mathbb{P} \text{Identifier}$$

$$\begin{aligned} &\forall i : \text{Instantiation}; TM : \mathbb{P} \text{TME} \bullet \\ &\quad \exists g : \text{Identifier} \mid g = \text{MatchGoal}(i.\text{match}, TM) \bullet \\ &\quad \text{InstantiationsLHSOperators}(i, TM) = \\ &\quad \quad \{o : \text{Identifier} \mid \text{makeWME}(g, \text{'OPERATOR'}, o, \text{No}) \in i.\text{match}.\text{lhs} \vee \\ &\quad \quad \text{makeWME}(g, \text{'OPERATOR'}, o, \text{Yes}) \in i.\text{match}.\text{lhs}\} \end{aligned}$$

An instantiation creates an operator if it makes an acceptable or require preference for an operator slot of a goal.

$$\text{InstantiationsRHSOperators} : \text{Instantiation} \times \mathbb{P} \text{TME} \rightarrow \mathbb{P} \text{Identifier}$$

$$\begin{aligned} &\forall i : \text{Instantiation}; TM : \mathbb{P} \text{TME} \bullet \\ &\quad \text{InstantiationsRHSOperators}(i, TM) = \\ &\quad \quad \{o : \text{Identifier} \mid \\ &\quad \quad \quad \text{makeUnaryPreference}(\text{MatchGoal}(i.\text{match}, TM), \\ &\quad \quad \quad \text{'OPERATOR'}, o, \text{'-'}) \in i.\text{rhs} \vee \\ &\quad \quad \quad \text{makeUnaryPreference}(\text{MatchGoal}(i.\text{match}, TM), \\ &\quad \quad \quad \text{'OPERATOR'}, o, \text{'!'}) \in i.\text{rhs}\} \end{aligned}$$

An instantiation only confers operator creation support if its match matched a working memory element whose identifier is in the transitive closure of the match goal's state.

$$\text{OperatorCreationLHS} : \mathbb{P} \text{TME} \rightarrow \mathbb{P} \text{Instantiation}$$

$$\begin{aligned} &\forall TM : \mathbb{P} \text{TME}; i : \text{Instantiation} \bullet \\ &\quad i \in \text{OperatorCreationLHS}(TM) \Leftrightarrow \\ &\quad \quad (\exists g : \{\text{MatchGoal}(i.\text{match}, TM)\} \bullet \\ &\quad \quad \quad (\exists s : \{\text{GoalState}(g, TM)\} \bullet \\ &\quad \quad \quad \quad (\text{OfId}_{\text{TME}}((\text{TCI}_{\text{TME}}(TM))(\{s\}), TM) \cap \\ &\quad \quad \quad \quad \quad \text{WME}_{\text{TME}}\{i.\text{match}.\text{lhs}\} \neq \emptyset))) \end{aligned}$$

An instantiation's right hand side only confers operator support on a preference if the preference's identifier is in the transitive closure of an instantiation's operator over temporary memory and the instantiation's right hand side.

$$\begin{array}{l}
 \text{OperatorCreationRHS} : \mathbb{P} TME \rightarrow (\text{Instantiation} \leftrightarrow \text{Preference}) \\
 \hline
 \forall TM : \mathbb{P} TME; i : \text{Instantiation}; p : \text{Preference} \bullet \\
 ((i, p) \in \text{OperatorCreationRHS}(TM)) \Leftrightarrow \\
 (\text{PreferencesId}(p) \in \\
 (\text{TCI}_{TME}(TM \cup \text{PreferenceTME}(i.\text{rhs}))) \\
 (\text{InstantiationsRHSOperators}(i, TM)))
 \end{array}$$

An instantiation confers operator creation support on a preference if its left hand side is an operator creation and the right hand side confers operator creation support on the preference.

$$\begin{array}{l}
 \text{OperatorCreation} : \mathbb{P} TME \rightarrow (\text{Instantiation} \leftrightarrow \text{Preference}) \\
 \hline
 \forall TM : \mathbb{P} TME; i : \text{Instantiation}; p : \text{Preference} \bullet \\
 (i, p) \in \text{OperatorCreation}(TM) \\
 \Leftrightarrow i \in \text{OperatorCreationLHS}(TM) \wedge \\
 (i, p) \in \text{OperatorCreationRHS}(TM)
 \end{array}$$

An instantiation can confer operator modification support if its left hand side matches an operator and its match matches a working memory element of the state.

$$\begin{array}{l}
 \text{OperatorModificationLHS} : \mathbb{P} TME \rightarrow \mathbb{P} \text{Instantiation} \\
 \hline
 \forall TM : \mathbb{P} TME; i : \text{Instantiation} \bullet \\
 i \in \text{OperatorModificationLHS}(TM) \Leftrightarrow \\
 (\exists g : \{\text{MatchGoal}(i.\text{match}, TM)\} \bullet \\
 (\exists s : \{\text{GoalState}(g, TM)\} \bullet \\
 ((\text{OfId}_{TME}(\text{TCI}_{TME}(TM))(\{s\}), TM) \cap \\
 \text{WME}_{TME}(i.\text{match}.\text{lhs}) \neq \emptyset) \wedge \\
 (\text{InstantiationsLHSOperators}(i, TM) \neq \emptyset)))
 \end{array}$$

An instantiation's right hand side confers operator modification support on a preference if its identifier is in the transitive closure of a matched operator.

$$\begin{array}{l}
 \text{OperatorModificationRHS} : \mathbb{P} \text{ TME} \rightarrow (\text{Instantiation} \leftrightarrow \text{Preference}) \\
 \hline
 \forall TM : \mathbb{P} \text{ TME}; i : \text{Instantiation}; p : \text{Preference} \bullet \\
 (i, p) \in \text{OperatorModificationRHS}(TM) \Leftrightarrow \\
 (\exists g : \{\text{MatchGoal}(i.\text{match}, TM)\}) \bullet \\
 \text{PreferencesId}(p) \in \\
 (\text{TCI}_{\text{TME}}(TM \cup \text{PreferenceTME}(i.\text{rhs}))) \\
 (\text{InstantiationsLHSOperators}(i, TM))
 \end{array}$$

An instantiation confers operator modification support on a preference if the left hand side is an operator modification left hand side and the right hand side confers operator modification support on the preference.

$$\begin{array}{l}
 \text{OperatorModification} : \mathbb{P} \text{ TME} \rightarrow (\text{Instantiation} \leftrightarrow \text{Preference}) \\
 \hline
 \forall TM : \mathbb{P} \text{ TME}; i : \text{Instantiation}; p : \text{Preference} \bullet \\
 (i, p) \in \text{OperatorModification}(TM) \Leftrightarrow \\
 i \in \text{OperatorModificationLHS}(TM) \wedge \\
 (i, p) \in \text{OperatorModificationRHS}(TM)
 \end{array}$$

An instantiation can confer operator application support if its left hand side matches its match goal's operator slot and a working memory element in the transitive closure of the match goal's state.

$$\begin{array}{l}
 \text{OperatorApplicationLHS} : \mathbb{P} \text{ TME} \rightarrow \mathbb{P} \text{ Instantiation} \\
 \hline
 \forall TM : \mathbb{P} \text{ TME}; i : \text{Instantiation} \bullet \\
 i \in \text{OperatorApplicationLHS}(TM) \Leftrightarrow \\
 (\exists g : \{\text{MatchGoal}(i.\text{match}, TM)\}) \bullet \\
 (\exists s : \{\text{GoalState}(g, TM)\}) \bullet \\
 (\exists o : \{\text{GoalOperator}(g, TM)\}) \bullet \\
 \text{makeWME}(g, \text{OPERATOR}, o, No) \in i.\text{match}.\text{lhs} \wedge \\
 (\text{OfId}_{\text{TME}}((\text{TCI}_{\text{TME}}(TM))(\{s\}), TM) \cap \\
 \text{WME}_{\text{TME}}(i.\text{match}.\text{lhs}) \neq \emptyset))
 \end{array}$$

An instantiation can confer operator application support on a preference if the preference's identifier is in the transitive closure of the state.

$$\begin{array}{l}
 \text{OperatorApplicationRHS} : \mathbb{P} TME \rightarrow (\text{Instantiation} \leftrightarrow \text{Preference}) \\
 \hline
 \forall TM : \mathbb{P} TME; i : \text{Instantiation}; p : \text{Preference} \bullet \\
 (i, p) \in \text{OperatorApplicationRHS}(TM) \Leftrightarrow \\
 (\exists g : \{\text{MatchGoal}(i.\text{match}, TM)\} \bullet \\
 (\exists s : \{\text{GoalState}(g, TM)\} \bullet \\
 \text{PreferencesId}(p) \in \text{TCI}_{TME}(TM \cup \text{PreferenceTME}(i.\text{rhs}))(\{s\})))
 \end{array}$$

An instantiation confers operator application support on a preference if the left hand side is an operator application left hand side and the right hand side confers operator application support on the preference.

$$\begin{array}{l}
 \text{OperatorApplication} : \mathbb{P} TME \rightarrow (\text{Instantiation} \leftrightarrow \text{Preference}) \\
 \hline
 \forall TM : \mathbb{P} TME; i : \text{Instantiation}; p : \text{Preference} \bullet \\
 (i, p) \in \text{OperatorApplication}(TM) \Leftrightarrow \\
 i \in \text{OperatorApplicationLHS}(TM) \wedge \\
 (i, p) \in \text{OperatorApplicationRHS}(TM)
 \end{array}$$

An instantiation confers operator support on a preference if the instantiation confers operator creation, modification or application support on the preference.

$$\begin{array}{l}
 \text{OSupport} : \mathbb{P} TME \rightarrow (\text{Instantiation} \leftrightarrow \text{Preference}) \\
 \hline
 \forall TM : \mathbb{P} TME; i : \text{Instantiation}; p : \text{Preference} \bullet \\
 (i, p) \in \text{OSupport}(TM) \Leftrightarrow \\
 (i, p) \in (\text{OperatorCreation}(TM) \cup \\
 \text{OperatorModification}(TM) \cup \\
 \text{OperatorApplication}(TM))
 \end{array}$$

5.9 Preference Phase Operations

The preference phase has eight states: an initial state, a final state, a state for retracting instantiations, a state for instantiating matches, a state for conferring support, a state for chunking an instantiation, a state for O-rejecting preferences and a state for changing the contents of preference memory based upon the contents of the support memories. This section specifies the operations of these states in seven sections.

5.9.1 Initialize, Finish, Reset and Quiescence

The preference phase waits in its final state until it is initialized with the *PPhaseInitialize* operation.

<i>PPhaseInitialize</i>
ΔRM
$pphase_state = PPhaseFinishedState$
$pphase_state' = PPhaseInitialState$

The *PPhaseReset* operation resets the preference phase state machine. The set of waiting O-reject preferences is emptied so that the next execution of the preference phase does not reject things that no longer have O-supported rejects.

<i>PPhaseReset</i>
ΔRM
$pphase_state' = PPhaseFinishedState$
$ORM' = \emptyset$

The top level state machine checks the state predicate, *Quiescence*, before it allows the preference phase to run. The system is quiescent if there is no instantiation waiting to be retracted, or match waiting to be instantiated.

<i>Quiescence</i>
RM
$\neg (\exists i : IS \bullet i.match \notin MS)$
$\neg (\exists m : MS \bullet \neg (\exists i : IS \bullet i.match = m))$

5.9.2 Retracting Instantiations

Preference phase begins its execution by moving from the initial state into a state that retracts instantiations.

<i>PPhaseStartRetractInstantiation</i>
ΔRM
$pphase_state = PPhaseInitialState$
$pphase_state' = PPhaseRetractInstantiationState$

When there exists an instantiation in instantiation memory whose match is no longer in the match set, then *PPhaseRetractInstantiation* removes the instantiation from the instantiation set. The instantiation support for each preference of the right hand side is removed from I-support memory.

$$\begin{array}{l}
 \text{PPhaseRetractInstantiation} \\
 \hline
 \Delta RM \\
 \hline
 pphase_state = PPhaseRetractInstantiationState = pphase_state' \\
 \exists i : IS \mid i.match \notin MS \bullet \\
 (IS' = IS \setminus \{i\} \wedge \\
 ISM' = ISM \setminus \{p : i.rhs \bullet (p \rightarrow i)\})
 \end{array}$$

5.9.3 Instantiate Match

When there is no instantiation in the instantiation set whose match is no longer satisfied, preference phase moves into the instantiate match state.

$$\begin{array}{l}
 \text{PPhaseStartInstantiateMatch} \\
 \hline
 \Delta RM \\
 \hline
 pphase_state = PPhaseRetractInstantiationState \\
 \neg (\exists i : IS \bullet i.match \notin MS) \\
 pphase_state' = PPhaseInstantiateMatchState
 \end{array}$$

PPhaseInstantiateMatch selects a match from the match set that does not yet have an instantiation and instantiates it. An instantiation is created and added to the instantiation memory, and recognition memory's *instantiation* variable is set to this new instantiation. Each newly generated preference is given a new instance numbering in *PM #*

PPhaseInstantiateMatch

ΔRM

pphase_state = *PPhaseInstantiateMatchState*

$\exists m : MS \mid \neg (\exists i : IS \bullet i.match = m) \bullet$
 $(\exists i : Instantiation; b : Binding \mid$
 $i.match = m \wedge$
 $i.lhs = m.production.lhs \wedge$
 $i.matching =$
 $InstantiateConditionMatching(m.matching,$
 $m.production.lhs \setminus \text{ran } PositiveCondition, m.binding) \wedge$
 $b ConsistentExtension m.binding \wedge$
 $b Covers (\bigcup (MakesComponents(m.production.rhs)) \cap Variable) \wedge$
 $i.binding = b \wedge$
 $i.rhs = (MakesPreference(b))(m.production.rhs) \bullet$
 $instantiation' = i \wedge$
 $IS' = IS \cup \{i\} \wedge$
 $\exists new_preferences : P Preference \mid new_preferences = (i.rhs \setminus PM) \bullet$
 $maz_preference' = maz_preference + \#new_preferences \bullet$
 $\exists f : new_preferences \mapsto ((maz_preference + 1) .. maz_preference') \bullet$
 $PM\#' = PM\# \oplus f)$

pphase_state' = *PPhaseConferSupportState*

5.9.4 Conferring Support

The *PPhaseConferSupport* operation adds support to the I-support memory for each right hand side preference, and if the preference is O-supported then it is given O-support for the match goal. Each reject preference that has O-support is added to the memory of O-reject preferences. After *PPhaseConferSupport* operates, chunking is stepped for the new instantiation.

<p><i>PPhaseConferSupport</i> _____</p> <p>ΔRM</p> <hr/> <p><i>pphase_state</i> = <i>PPhaseConferSupportState</i></p> <p>$ISM' = ISM \cup \{p : instantiation.rhs \bullet (p, instantiation)\} \wedge$ $OSM' = OSM \cup$ $\{p : instantiation.rhs \mid$ $(instantiation, p) \in OSupport(TM) \bullet$ $(p, MatchGoal(instantiation.match, TM))\} \wedge$ $ORM' = ORM \cup$ $\{p : instantiation.rhs \cap ran UnaryP \mid$ $(UnaryP \sim (p)).preference = \text{reject} \wedge (instantiation, p) \in OSupport(TM)\}$</p> <p><i>ChunkingInitialize</i></p> <p><i>pphase_state'</i> = <i>PPhaseChunkingState</i></p>

5.9.5 Stepping Chunking

PPhaseStepChunking steps the chunking state machine through chunking the instantiation.

<p><i>PPhaseStepChunking</i> _____</p> <p>ΔRM</p> <hr/> <p><i>pphase_state</i> = <i>PPhaseChunkingState</i> = <i>pphase_state'</i></p> <p><i>ChunkingStep</i></p>
--

Chunking may return a new instantiation or it may produce no new instantiation if the input instantiation is for the top goal. The variable *chunking_result* signals the existence of a new instantiation to preference phase.

If there is a new instantiation, it must have its support calculated, so preference phase returns to the confer support state. Once the support has been conferred, the new instantiation will be sent back into chunking so that its effects may be chunked.

<i>PPhaseChunkingReturnedInstantiation</i>
ΔRM
<i>pphase_state</i> = <i>PPhaseChunkingState</i>
- pre <i>Chunking</i>
<i>chunking_result</i> = <i>NewInstantiation</i>
<i>pphase_state'</i> = <i>PPhaseConferSupportState</i>

If no instantiation has been generated by chunking, the preference phase returns to the instantiate match state to pick a new match for instantiation.

<i>PPhaseChunkingNoInstantiation</i>
ΔRM
<i>pphase_state</i> = <i>PPhaseChunkingState</i>
- pre <i>Chunking</i>
<i>chunking_result</i> = <i>NoInstantiation</i>
<i>pphase_state'</i> = <i>PPhaseInstantiateMatchState</i>

5.9.6 O-rejects

When there are no more matches to instantiate, preference phase starts the processing of O-rejects.

<i>PPhaseStartOReject</i>
ΔRM
<i>pphase_state</i> = <i>PPhaseInstantiateMatchState</i>
$\forall m : MS \bullet \exists i : IS \bullet i.match = m$
<i>pphase_state'</i> = <i>PPhaseORejectState</i>

PPhaseORejectPreferences chooses an O-reject preference from the O-reject memory and deletes it from memory. It also deletes from preference memory all preferences that share the same identifier, attribute and value. The I-support and O-support memories are also cleaned of all support for the rejected preferences, including the O-reject itself. The \Leftarrow operation ([Spi89] 99) restricts a relation or partial function to apply only to those elements of its domain that are not in the provided set.

$$\begin{array}{l}
 \text{--- } PPhaseORejectPreferences \text{ ---} \\
 \Delta RM \\
 \text{---} \\
 pphase_state = PPhaseORejectState = pphase_state' \\
 \exists r : ORM; ur : UnaryPreference \quad ur = UnaryP^{\sim}(r) \bullet \\
 (ORM' = ORM \setminus \{r\} \wedge \\
 (\exists P : \mathcal{P} Preference \\
 P = UnaryP(\{up : UnaryP^{\sim}(((PM \cup dom ISM) \cap ran UnaryP)\}) \\
 up.id = ur.id \wedge up.attribute = ur.attribute \wedge up.value = ur.value\}) \cup \\
 BinaryP(\{bp : BinaryP^{\sim}(((PM \cup dom ISM) \cap ran BinaryP)\}) \\
 bp.id = ur.id \wedge bp.attribute = ur.attribute \wedge bp.value = ur.value\}) \cup \\
 \{r\} \bullet \\
 PM' = PM \setminus P \wedge \\
 ISM' = P \Leftarrow ISM \wedge \\
 OSM' = P \Leftarrow OSM))
 \end{array}$$

5.9.7 Changing Preference Memory

When all of the O-rejects have been processed, preference phase starts changing preference memory to correspond with the contents of the support memories.

$$\begin{array}{l}
 \text{--- } PPhaseStartChangePM \text{ ---} \\
 \Delta RM \\
 \text{---} \\
 pphase_state = PPhaseORejectState \\
 ORM = \emptyset \\
 pphase_state' = PPhaseChangePMState
 \end{array}$$

If there is a preference that has support in I-support memory or O-support memory that is not in preference memory, then it is added to preference memory, and its slot is added to the set of changed slots.

$$\begin{array}{l}
 \text{--- } PPhaseAddPreference \text{ ---} \\
 \Delta RM \\
 \hline
 pphase_state = PPhaseChangePMState = pphase_state' \\
 \exists p : \text{dom } ISM \cup \text{dom } OSM \mid p \notin PM \bullet \\
 (PM' = PM \cup \{p\} \wedge \\
 \text{changed_slots}' = \text{changed_slots} \cup \{(p.id, p.attribute)\})
 \end{array}$$

If there is a preference in preference memory that does not have I-support or O-support, it is removed from preference memory, and its slot is added to the changed slots.

$$\begin{array}{l}
 \text{--- } PPhaseRemovePreference \text{ ---} \\
 \Delta RM \\
 \hline
 pphase_state = PPhaseChangePMState = pphase_state' \\
 \exists p : PM \mid p \notin \text{dom } ISM \cup \text{dom } OSM \bullet \\
 (PM' = PM \setminus \{p\} \wedge \\
 \text{changed_slots}' = \text{changed_slots} \cup \{(p.id, p.attribute)\})
 \end{array}$$

When there are no more differences between the supported preferences and the contents of preference memory, preference phase is complete.

$$\begin{array}{l}
 \text{--- } PPhaseFinish \text{ ---} \\
 \Delta RM \\
 \hline
 pphase_state = PPhaseChangePMState \\
 - (\exists p : PM \bullet p \notin \text{dom } ISM \cup \text{dom } OSM) \\
 - (\exists p : \text{dom } ISM \bullet p \notin PM) \\
 pphase_state' = PPhaseFinishedState
 \end{array}$$

PPhaseStep moves the preference phase state machine through one of its possible transitions.

$$\begin{aligned} PPhaseStep \hat{=} & \\ & PPhaseStartRetractInstantiation \vee PPhaseRetractInstantiation \vee \\ & PPhaseStartInstantiateMatch \vee PPhaseInstantiateMatch \vee \\ & PPhaseStartOReject \vee PPhaseORejectPreferences \vee \\ & PPhaseStartChangePM \vee PPhaseAddPreference \vee \\ & PPhaseRemovePreference \vee PPhaseFinish \end{aligned}$$

Chapter 6

IO

This chapter specifies Soar's mechanism for communicating with the external world, called SoarIO or IO for short. IO's operation is graphically represented in Figure 6.1. IO's two state machines, *InputCycle* and *OutputCycle*, are stepped by the top level state machine. The *InputCycle* reads *InputStructures* from *InputChannels* and places their results on *InputAttributes* on the top state, and in the transitive closure of the value of the input attribute. The *OutputCycle* packages the transitive closures of output attributes of the top state into *OutputStructures* and ships them out *OutputChannels*.

6.1 Channels, Attributes and IO Mappings

The user of Soar provides IO with two fixed sets: *InputChannel* and *OutputChannel*. For the purposes of the specification we represent these channels as a sets of constant names for the channels. The implementation may choose to represent a channel as a function name, or a pointer to a function; in essence it is a computational operation that provides input perceptions or receives motor operations.

The user also specifies two sets of *InputAttributes* and *OutputAttributes*. The *InputAttributes* are the attributes that input channels use to augment the top level state with their perceptions. The user provided *InputMapping* maps each *InputChannel* to the *InputAttribute* that receives its perceptions. Similarly, the *OutputAttributes* are the attributes of the top level state from which IO collects their transitive closure and ships it to the *OutputChannel* specified by the *OutputMapping*.

No channel can be both an input and an output channel, and no attribute can be both an input and an output attribute.

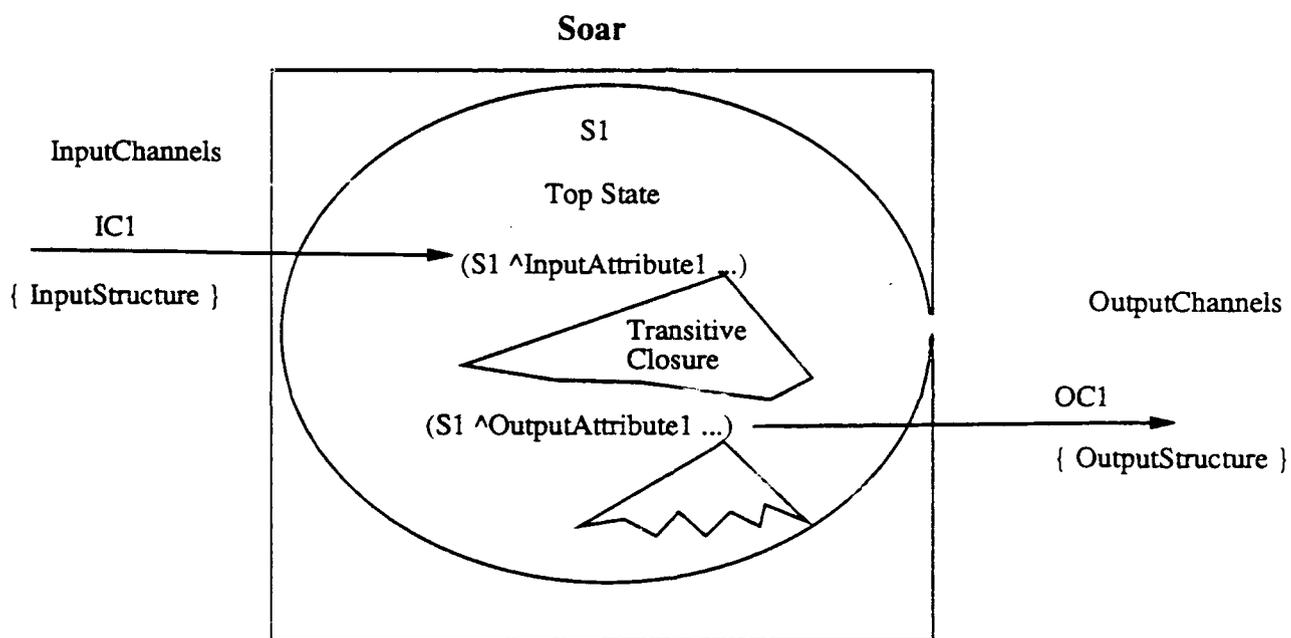


Figure 6.1: Graphical Representation of IO

InputChannel : P Constant
InputAttribute : P Constant
OutputChannel : P Constant
OutputAttribute : P Constant

disjoint (*InputChannel*, *OutputChannel*)
 disjoint (*InputAttribute*, *OutputAttribute*)

InputMapping : *InputChannel* → *InputAttribute*
OutputMapping : *OutputAttribute* → *OutputChannel*

6.2 Output Structures

The *OutputCycle* collects output into packages called *OutputStructures*. On each cycle all *OutputAttributes* have the working memory elements in their transitive closure collected into an *OutputStructure* and shipped out their corresponding *OutputChannel*. The predicate of the output structure ensures that the set of working memory elements is transitively closed from the object identifier.

OutputStructure
object : Identifier
wmes : P WME
 $OfId_{WME}((TCI_{WME}(wmes))(\{object\}), wmes) = wmes$

6.3 Input Structures

On each *InputCycle*, all of the *InputChannels* are polled to produce changes to working memory. These changes are bundled into *InputStructures*. There are three types of input structures: a *NewInputStructure*, a *ModifyInputStructure* and a *DeleteInputStructure*.

A *NewInputStructure* tells Soar to augment the top state with a working memory of value *new_object* and the attribute of the input channel, and to also add all of the elements in *wmes*.

NewInputStructure
new_object : Identifier
wmes : P WME
 $OfId_{WME}((TCI_{WME}(wmes))(\{new_object\}), wmes) = wmes$

A *ModifyInputStructure* tells Soar to find the augmentation of the top state that holds the input channel's attribute and the *modify_object* value. Soar must add all of the elements in the *adds* set, and delete all of the elements in the *deletes* set.

<i>ModifyInputStructure</i> <i>modify_object</i> : Identifier <i>adds, deletes</i> : P WME
--

The *DeleteInputStructure* tells Soar to delete the augmentation of the input channel's attribute and the *delete_object* value and all of the elements in its transitive closure.

<i>DeleteInputStructure</i> <i>delete_object</i> : Identifier
--

The free type *InputStructure* is used to construct a disjunctive type of the basic input types.

```

InputStructure ::= New<<NewInputStructure>>
                | Modify<<ModifyInputStructure>>
                | Delete<<DeleteInputStructure>>

```

6.4 Cycles of IO

As Soar behaves, it reads in parallel from all of its input channels. Each input channel is free to send more than one *InputStructure* to the system. A *CycleOfInput* is defined to capture the input for a single cycle. It maps each *InputChannel* to the set of inputs read during that cycle.

```

CycleOfInput == InputChannel → P InputStructure
CycleOfOutput == OutputChannel → OutputStructure

```

Similarly, each parallel ply of Soar's motor actions are captured in a partial function from *OutputChannels* to the single possible *OutputStructure* that Soar produced.

6.5 IO State

IO requires two state machines: one for the *InputCycle* (Figure 6.2) and one for the *OutputCycle* (Figure 6.3).

The *InputCycle* requires three states.

1. *InputCycleInitialState* — the initial state of the machine
2. *InputCycleReadState* — the state in which a single *InputStructure* is read from an *InputChannel*
3. *InputCycleFinishedState* — the final state of the machine

```
InputCycleState ::= InputCycleInitialState  
                  | InputCycleReadState  
                  | InputCycleFinishedState
```

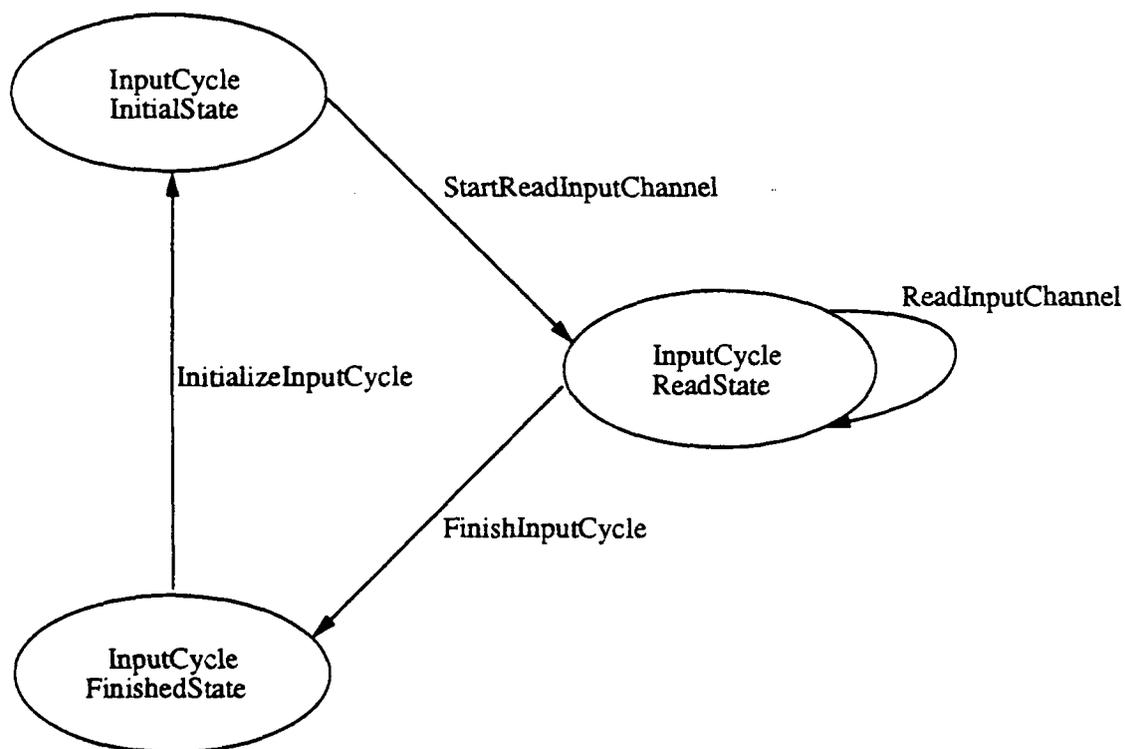


Figure 6.2: Input Cycle State Machine

The *OutputCycle* also requires only three states:

1. *InputCycleInitialState* — the initial state of the machine.
2. *InputCycleReadState* — the state in which an *OutputStructure* is collected from the top level state and sent out an *OutputChannel*.
3. *InputCycleFinishedState* — the final state of the machine.

```
OutputCycleState ::= OutputCycleInitialState  
                   | OutputCycleSendState  
                   | OutputCycleFinishedState
```

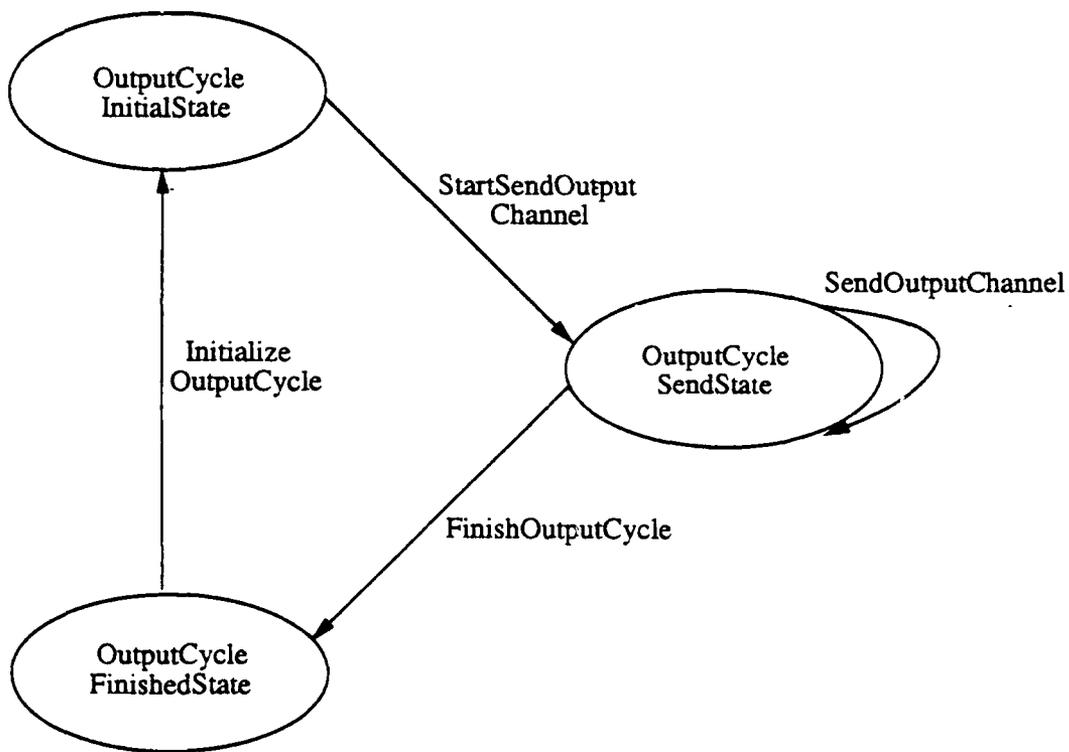


Figure 6.3: Output Cycle State Machine

The *IO* state schema shares working memory and goal memory with the other major modules of Soar. The *IO* system searches working memory to keep updated the top state of the goal stack. If there is no top state, then *TopState* is set to a special symbol 'NIL'.

The schema holds state counters *input_cycle_state* and *output_cycle_state* for the two IO state machines.

The specification of *IO* takes a traditional approach to modeling input/output behavior by using sequences of inputs to model the sequence of inputs that Soar would encounter over time, and by recording its outputs over time in a sequence.

inputs is the sequence of *CycleOfInput* that models a possibly infinite stream of cycles of input that Soar will perceive as it behaves. *outputs* is the initially empty sequence of *CycleOfOutput* in which IO records Soar's output over time.

Other approaches to modeling the input/output behavior of a computational system are possible. For example, the operations that read input from the sequence of inputs can be replaced by an under-specified operation that produces any one input from the set of all possible inputs. Similarly, the output operation can be under-specified to have no effect.

<i>IO</i>
$WM : \mathcal{P} WME$ $GM : \mathcal{P} Identifier$ $TopState : Identifier \cup \{ 'NIL' \}$ $inputs : seq CycleOfInput$ $input_cycle_state : InputCycleState$ $outputs : seq CycleOfOutput$ $output_cycle_state : OutputCycleState$
$\exists g : GM \mid makeWME(g, 'OBJECT', 'NIL') \in WM \bullet$ $((\exists s : Identifier \bullet makeWME(g, 'STATE', s) \in WM) \Rightarrow$ $(\exists s : Identifier \mid makeWME(g, 'STATE', s) \in WM \bullet$ $TopState = S)$ $\neg (\exists s : Identifier \bullet makeWME(g, 'STATE', s) \in WM) \Rightarrow$ $TopState = 'NIL')$

IO starts with both state machines in their finished state and an empty sequence of outputs.

<i>InitIO</i>
IO $input_cycle_state = InputCycleFinishedState$ $output_cycle_state = OutputCycleFinishedState$ $outputs = \langle \rangle$

6.6 The Input Cycle

The *InputCycle* is defined using five transitions.

1. *InputCycleInitialize* — this transition initializes the state machine.
2. *StartReadInputChannel* — this transition starts the processing of input.
3. *ReadInputChannel* — this transition is a very complex composite transition. With each step, it reads one part of any of the three types of *InputStructures*.
4. *InputCycleFinish* — this transition finishes the *InputCycle* when all of the input has been read.
5. *InputCycleReset* — this transition resets the state machine.

6.6.1 Legal Input

The *InputCycle* is the only place in the specification of Soar that an error condition is checked. An example of an errorful input is an input structure that attempts to remove working memory elements that are not in the transitive closure of its input attribute.

The specification can check that its input sequence is correct before execution, but as the implementation does not have the full sequence of inputs available at start up time, it must check them one by one. For consistency we have specified that each cycle of input is checked for correctness as they are popped from the sequence of inputs. This subsection defines the legal input predicate to test each input.

An identifier is an *InputLink* of an attribute, if it there is an augmentation of the top state in working memory that holds it as a value.

$$\begin{array}{l}
 \text{InputLink}_- : \mathbb{P} \text{ WME} \times \text{Identifier} \times \text{InputAttribute} \rightarrow \text{Identifier} \\
 \forall \text{ WM} : \mathbb{P} \text{ WME}; s : \text{Identifier}; \text{InputAttribute} : \text{InputAttribute}; \\
 id : \text{Identifier}; link : \text{Preference} \bullet \\
 \text{InputLink}((\text{WM}, s, \text{InputAttribute}), id) \Leftrightarrow \\
 (\exists w : \text{WM} \bullet w.id = s \wedge w.attribute = \text{InputAttribute} \wedge \\
 w.value = id \wedge w.context_acceptable_preference = \text{No})
 \end{array}$$

LegalInput checks the first *CycleOfInput* in the *inputs* sequence for several properties:

- that no new and modify are for the same object
- that no new and delete are for the same object
- that no delete and modify are for the same object
- that no two news are for the same object
- that no two modifies are for the same object

- that the adds of different modifies are disjoint, and that one modify is not adding what another is deleting
- that a *NewInputStructure* is for an *InputLink* that does not exist on the top state
- that a *DeleteInputStructure* is for an existing *InputLink* and
- that a *ModifyInputStructure* is for an existing *InputLink*, and that its adds and deletes modify the transitive closure of the input object.

— *LegalInput* —
IO

$$\begin{aligned}
& \forall \text{channel} : \text{InputChannel} \bullet \\
& \quad \forall \text{is1}, \text{is2} : \text{inputs}(1)(\text{channel}) \bullet \\
& \quad \quad \text{is1} \in \text{ran New} \wedge \text{is2} \in \text{ran New} \Rightarrow \\
& \quad \quad \quad ((\text{New}^{\sim}(\text{is1})).\text{new_object} = (\text{New}^{\sim}(\text{is2})).\text{new_object} \Rightarrow \text{is1} = \text{is2}) \wedge \\
& \quad \quad \text{is1} \in \text{ran Modify} \wedge \text{is2} \in \text{ran Modify} \Rightarrow \\
& \quad \quad \quad ((\text{Modify}^{\sim}(\text{is1})).\text{modify_object} = \\
& \quad \quad \quad \quad (\text{Modify}^{\sim}(\text{is2})).\text{modify_object} \Rightarrow \text{is1} = \text{is2} \wedge \\
& \quad \quad \quad \text{disjoint} \langle (\text{Modify}^{\sim}(\text{is1})).\text{adds}, (\text{Modify}^{\sim}(\text{is2})).\text{adds} \rangle \wedge \\
& \quad \quad \quad \text{disjoint} \langle (\text{Modify}^{\sim}(\text{is1})).\text{adds}, (\text{Modify}^{\sim}(\text{is2})).\text{deletes} \rangle) \wedge \\
& \quad \quad \text{is1} \in \text{ran New} \wedge \text{is2} \in \text{ran Modify} \Rightarrow \\
& \quad \quad \quad ((\text{New}^{\sim}(\text{is1})).\text{new_object} \neq (\text{Modify}^{\sim}(\text{is2})).\text{modify_object}) \wedge \\
& \quad \quad \text{is1} \in \text{ran New} \wedge \text{is2} \in \text{ran Delete} \Rightarrow \\
& \quad \quad \quad ((\text{New}^{\sim}(\text{is1})).\text{new_object} \neq (\text{Delete}^{\sim}(\text{is2})).\text{delete_object}) \wedge \\
& \quad \quad \text{is1} \in \text{ran Delete} \wedge \text{is2} \in \text{ran Modify} \Rightarrow \\
& \quad \quad \quad ((\text{Delete}^{\sim}(\text{is2})).\text{delete_object} \neq (\text{Modify}^{\sim}(\text{is2})).\text{modify_object}) \\
& \quad \forall \text{channel} : \text{InputChannel} \bullet \forall \text{is} : \text{inputs}(1)(\text{channel}) \bullet \\
& \quad \quad \text{is} \in \text{ran New} \Rightarrow \\
& \quad \quad \quad (\neg \text{InputLink}((\text{WM}, \text{TopState}, \text{InputMapping}(\text{channel})), \\
& \quad \quad \quad \quad (\text{New}^{\sim}(\text{is})).\text{new_object})) \wedge \\
& \quad \quad \text{is} \in \text{ran Delete} \Rightarrow \\
& \quad \quad \quad \text{InputLink}((\text{WM}, \text{TopState}, \text{InputMapping}(\text{channel})), \\
& \quad \quad \quad \quad (\text{Delete}^{\sim}(\text{is})).\text{delete_object}) \wedge \\
& \quad \quad \text{is} \in \text{ran Modify} \Rightarrow \\
& \quad \quad \quad (\text{InputLink}((\text{WM}, \text{TopState}, \text{InputMapping}(\text{channel})), \\
& \quad \quad \quad \quad (\text{Modify}^{\sim}(\text{is})).\text{modify_object}) \wedge \\
& \quad \quad \quad ((\text{Modify}^{\sim}(\text{is})).\text{deletes} \\
& \quad \quad \quad \quad \subseteq \text{OfId}_{\text{WME}}((\text{TCI}_{\text{WME}}(\text{WM}))(\{(\text{Modify}^{\sim}(\text{is})).\text{modify_object}\}), \text{WM}) \wedge \\
& \quad \quad \quad ((\text{Modify}^{\sim}(\text{is})).\text{adds} \\
& \quad \quad \quad \quad \subseteq \text{OfId}_{\text{WME}}((\text{TCI}_{\text{WME}}(\text{WM} \cup (\text{Modify}^{\sim}(\text{is})).\text{adds})) \\
& \quad \quad \quad \quad \quad (\{(\text{Modify}^{\sim}(\text{is})).\text{modify_object}\}, \text{WM} \cup (\text{Modify}^{\sim}(\text{is})).\text{adds}))))))
\end{aligned}$$

Although the specification defines the conditions of acceptable input, it does not specify how the implementation should handle unacceptable input.

6.6.2 Initializing and Starting the *InputCycle*

InputCycleInitialize moves the *InputCycle* state machine into its initial state.

$$\frac{\text{--- } \textit{InputCycleInitialize} \text{ ---}}{\Delta IO}$$

$$\textit{input_cycle_state} = \textit{InputCycleFinishedState}$$

$$\textit{input_cycle_state}' = \textit{InputCycleInitialState}$$

StartReadInput moves the machine into the *InputCycleReadState* to begin reading input. The whole *CycleOfInput* is checked for legality using the *LegalInput* predicate.

$$\frac{\text{--- } \textit{StartReadInput} \text{ ---}}{\Delta IO}$$

$$\textit{LegalInput}$$

$$\textit{input_cycle_state} := \textit{InputCycleInitialState}$$

$$\textit{input_cycle_state}' = \textit{InputCycleReadState}$$

6.6.3 Reading NewInputStructures

When the first *CycleOfInput* holds a *NewInputStructure*, *ReadNewInputStructureObject* reads its object. Chunking's *TraceWME* operation is called to give the working memory element a new instance number.

$$\frac{\text{--- } \textit{ReadNewInputStructureObject} \text{ ---}}{\Delta IO}$$

$$\textit{input_cycle_state} = \textit{InputCycleReadState} = \textit{input_cycle_state}'$$

$$\exists ic : \textit{InputChannel} \cdot ic \in \text{dom}(\textit{inputs}(1)) \bullet$$

$$\exists ia : \textit{InputAttribute} \cdot ia = \textit{InputMapping}(ic) \bullet$$

$$\exists n : \textit{inputs}(1)(ic) \cap \text{ran } \textit{New}; w : \textit{WME} :$$

$$\textit{makeWME}(\textit{TopState}, ic, (\textit{New} \sim (n)).\textit{new_object}, \textit{No}) \in \textit{WM} \wedge$$

$$w = \textit{makeWME}(\textit{TopState}, ic, (\textit{New} \sim (n)).\textit{new_object}, \textit{No}) \bullet$$

$$(\textit{WM}' = \textit{WM} \cup \{w\} \wedge$$

$$\textit{TraceWME})$$

After the top state has been augmented with the object, *ReadNewInputStructureWME* reads in one of the new working memory elements, and destructively removes it from the *NewInputStructure*.

<p style="text-align: center;"><i>ReadNewInputStructureWME</i></p> <hr/> <p style="text-align: center;">ΔIO</p> <hr/> <p><i>input_cycle_state</i> = <i>InputCycleReadState</i> = <i>input_cycle_state'</i></p> <p>$\exists ic : InputChannel \mid ic \in \text{dom}(inputs(1)) \bullet$</p> <p style="padding-left: 20px;">$\exists ia : InputAttribute \mid ia = InputMapping(ic) \bullet$</p> <p style="padding-left: 20px;">$\exists n : inputs(1)(ic) \cap \text{ran } New \bullet$</p> <p style="padding-left: 40px;">$\exists N : NewInputStructure \mid N = New \sim (n) \wedge$ $makeWME(TopState, ia, N.new_object, No) \in WM \bullet$</p> <p style="padding-left: 20px;">$\exists w : N.wmes \mid w \notin WM \bullet$</p> <p style="padding-left: 20px;">$\exists N' : NewInputStructure :$</p> <p style="padding-left: 40px;">$N'.new_object = N.new_object \wedge$ $N'.wmes = N.wmes \setminus \{w\} \bullet$ $(WM' = WM \cup \{w\} \wedge$ $TraceWME \wedge$ $inputs'(1) = inputs(1) \oplus$ $\{ic \mapsto (inputs(1)(ic) \setminus \{n\}) \cup \{New(N')\}\} \wedge$ $tail(inputs') = tail(inputs))$</p> <hr/>

When all of the working memory elements have been read in, *ReadNewInputStructureDone* pops the input from the first cycle of input.

<p style="text-align: center;"><i>ReadNewInputStructureDone</i></p> <hr/> <p style="text-align: center;">ΔIO</p> <hr/> <p><i>input_cycle_state</i> = <i>InputCycleReadState</i> = <i>input_cycle_state'</i></p> <p>$\exists ic : InputChannel \mid ic \in \text{dom}(inputs(1)) \bullet$</p> <p style="padding-left: 20px;">$\exists ia : InputAttribute \mid ia = InputMapping(ic) \bullet$</p> <p style="padding-left: 20px;">$\exists n : inputs(1)(ic) \cap \text{ran } New \bullet$</p> <p style="padding-left: 40px;">$\exists N : NewInputStructure \mid N = New \sim (n) \wedge$ $makeWME(TopState, ia, N.new_object, No) \in WM \bullet$</p> <p style="padding-left: 20px;">$((N.wmes = \emptyset \wedge$ $inputs'(1) = inputs(1) \oplus \{ic \mapsto inputs(1)(ic) \setminus \{n\}\} \wedge$ $tail(inputs') = tail(inputs))$</p> <hr/>

ReadNewInputStructure is the composite operation for reading an entire *NewInputStructure*.

$$\begin{aligned} \text{ReadNewInputStructure} \hat{=} & \text{ReadNewInputStructureObject} \vee \\ & \text{ReadNewInputStructureWME} \vee \\ & \text{ReadNewInputStructureDone} \end{aligned}$$

6.6.4 Reading ModifyInputStructures

When there is a *ModifyInputStructure* in the first *CycleOfInput*, *ReadModifyInputStructureAdd* reads in one of the elements that it should add, and deletes the element from the set of adds.

$$\begin{array}{l} \text{ReadModifyInputStructureAdd} \\ \hline \Delta IO \\ \hline \text{input_cycle_state} = \text{inputCycleReadState} = \text{input_cycle_state}' \\ \exists ic : \text{InputCycle} \mid ic \in \text{dom}(\text{inputs}(1)) \bullet \\ \quad \exists ia : \text{InputAttribute} \mid ia = \text{InputMapping}(ic) \bullet \\ \quad \exists m : \text{inputs}(1)(ic) \cap \text{ran New} \bullet \\ \quad \exists M : \text{ModifyInputStructure} \mid M = \text{Modify}^{\sim}(m) \wedge \\ \quad \quad \text{makeWME}(\text{TopState}, ia, M.\text{modify_object}, No) \in WM \bullet \\ \quad \exists w : M.\text{adds} \mid w \notin WM \bullet \\ \quad \exists M' : \text{ModifyInputStructure} \mid \\ \quad \quad M'.\text{modify_object} = M.\text{modify_object} \wedge \\ \quad \quad M'.\text{adds} = M.\text{adds} \setminus \{w\} \wedge \\ \quad \quad M'.\text{deletes} = M.\text{deletes} \bullet \\ \quad (WM' = WM \cup \{w\} \wedge \\ \quad \text{TraceWME} \wedge \\ \quad \text{inputs}'(1) = \text{inputs}(1) \oplus \\ \quad \quad \{ic \mapsto (\text{inputs}(1)(ic) \setminus \{m\}) \cup \{\text{Modify}(M')\}\} \wedge \\ \quad \text{tail}(\text{inputs}') = \text{tail}(\text{inputs})) \end{array}$$

ReadModifyInputStructureDelete reads in a delete of a modify, and removes the element from the modify's delete set.

$$\begin{array}{l}
 \text{--- } \textit{ReadModifyInputStructureDelete} \text{ ---} \\
 \Delta IO \\
 \text{---} \\
 \textit{input_cycle_state} = \textit{InputCycleReadState} = \textit{input_cycle_state}' \\
 \exists ic : \textit{InputChannel} \mid ic \in \text{dom}(\textit{inputs}(1)) \bullet \\
 \quad \exists ia : \textit{InputAttribute} \mid ia = \textit{InputMapping}(ic) \bullet \\
 \quad \quad \exists m : \textit{inputs}(1)(ic) \cap \text{ran } \textit{New} \bullet \\
 \quad \quad \quad \exists M : \textit{ModifyInputStructure} \mid M = \textit{Modify}^{\sim}(m) \wedge \\
 \quad \quad \quad \quad \textit{makeWME}(\textit{TopState}, ia, M.\textit{modify_object}, \textit{No}) \in \textit{WM} \bullet \\
 \quad \quad \quad \exists w : M.\textit{deletes} \mid w \in \textit{WM} \bullet \\
 \quad \quad \exists M' : \textit{ModifyInputStructure} \mid \\
 \quad \quad \quad M'.\textit{modify_object} = M.\textit{modify_object} \wedge \\
 \quad \quad \quad M'.\textit{deletes} = M.\textit{deletes} \setminus \{w\} \wedge \\
 \quad \quad \quad M'.\textit{adds} = M.\textit{adds} \bullet \\
 \quad \quad (\textit{WM}' = \textit{WM} \setminus \{w\} \wedge \\
 \quad \quad \textit{inputs}'(1) = \textit{inputs}(1) \oplus \\
 \quad \quad \quad \{ic \rightarrow \textit{inputs}(1)(ic) \setminus \{m\}\} \wedge \\
 \quad \quad \textit{tail}(\textit{inputs}') = \textit{tail}(\textit{inputs})) \\
 \text{---}
 \end{array}$$

When all of the changes have been processed, *ReadModifyInputStructureDone* pops the modify from the first cycle of inputs.

$$\begin{array}{l}
 \text{--- } \textit{ReadModifyInputStructureDone} \text{ ---} \\
 \Delta IO \\
 \text{---} \\
 \textit{input_cycle_state} = \textit{InputCycleReadState} = \textit{input_cycle_state}' \\
 \exists ic : \textit{InputChannel} \mid ic \in \text{dom}(\textit{inputs}(1)) \bullet \\
 \quad \exists ia : \textit{InputAttribute} \mid ia = \textit{InputMapping}(ic) \bullet \\
 \quad \quad \exists m : \textit{inputs}(1)(ic) \cap \text{ran } \textit{New} \bullet \\
 \quad \quad \exists M : \textit{ModifyInputStructure} \mid \\
 \quad \quad \quad M = \textit{Modify}^{\sim}(m) \wedge M.\textit{adds} = M.\textit{deletes} = \emptyset \bullet \\
 \quad \quad (\textit{inputs}'(1) = \textit{inputs}(1) \oplus \{ic \rightarrow \textit{inputs}(1)(ic) \setminus \{m\}\} \wedge \\
 \quad \quad \textit{tail}(\textit{inputs}') = \textit{tail}(\textit{inputs})) \\
 \text{---}
 \end{array}$$

ReadModifyInputStructure is the composite operation that joins all of the operations for reading a *ModifyInputStructure*.

$$\begin{array}{l}
 \textit{ReadModifyInputStructure} \hat{=} \textit{ReadModifyInputStructureAdd} \vee \\
 \quad \textit{ReadModifyInputStructureDelete} \vee \\
 \quad \textit{ReadModifyInputStructureDone}
 \end{array}$$

6.6.5 Reading DeleteInputStructures

ReadDeleteInputStructure pops a *DeleteInputStructure* from the first *CycleOffInput*, and removes the associated *InputLink*. The working memory elements in the transitive closure of the link are not explicitly deleted. The working memory elements that are no longer accessible from the context stack will eventually be removed by *decide*.

ReadDeleteInputStructure

ΔIO

$input_cycle_state = InputCycleReadState = input_cycle_state'$

$\exists ic : InputChannel \mid ic \in \text{dom}(inputs(1)) \bullet$

$\exists ia : InputAttribute \mid ia = InputMapping(ic) \bullet$

$\exists d . inputs(1)(ic) \cap \text{ran } Delete \bullet$

$(WM' = WM \setminus$

$\{makeWME(TopState, ia, (Delete \sim (d)).delete_object, No)\} \wedge$

$inputs'(1) = inputs(1) \oplus$

$\{ic \rightarrow inputs(1)(ic) \setminus \{d\}\} \wedge$

$tail(inputs') = tail(inputs))$

6.6.6 Reading an Input Channel

CloseInputChannel will close any channel that has no remaining input structures.

CloseInputChannel

ΔIO

$input_cycle_state = InputCycleReadState = input_cycle_state'$

$TopState \neq \text{'NIL'}$

$\exists ic : InputChannel \mid ic \in \text{dom}(inputs(1)) \bullet$

$(inputs(1)(ic) = \emptyset \wedge$

$inputs'(1) = \{ic\} \triangleleft inputs(1) \wedge$

$tail(inputs') = tail(inputs))$

ReadInputChannelAnyInputStructure will take one step of reading any input structure of the first cycle of input.

$$\begin{aligned} \textit{ReadInputChannel} \hat{=} & \textit{ReadNewInputStructure} \vee \\ & \textit{ReadDeleteInputStructure} \vee \\ & \textit{ReadModifyInputStructure} \vee \\ & \textit{CloseInputChannel} \end{aligned}$$

6.6.7 Finishing, Stepping and Resetting the *InputCycle*

When there are no more input structures to read on any input channels, *InputCycleFinish* pops the inputs vector and moves the input cycle to the finished state.

$$\begin{array}{l} \textit{InputCycleFinish} \\ \hline \Delta IO \\ \textit{input_cycle_state} = \textit{InputCycleFinishedState} \\ \textit{pre ReadInputChannel} \\ \textit{inputs}' = \textit{tail}(\textit{inputs}) \\ \textit{input_cycle_state}' = \textit{InputCycleFinishedState} \end{array}$$

InputCycleStep steps the input cycle using the three operations.

$$\textit{InputCycleStep} \hat{=} \textit{StartReadInput} \vee \textit{ReadInputChannel} \vee \textit{InputCycleFinish}$$

InputCycleReset is specified to allow the implementation to reset the operation of the input cycle state machine at any point of execution.

$$\begin{array}{l} \textit{InputCycleReset} \\ \hline \Delta IO \\ \textit{input_cycle_state}' = \textit{InputCycleInitialState} \end{array}$$

6.7 The Output Cycle

The *OutputCycle* state machine requires only five transitions.

1. *OutputChannelInitialize* — this operation initializes the *OutputCycle* machine.
2. *StartSendOutputChannel* — this operation moves the *OutputCycle* from the initial state into the state that sends out output.
3. *SendOutputChannel* — this composite operation ships an *OutputStructure* out an *OutputChannel*.
4. *OutputCycleFinish* — when there are no more outputs to send, this operation ends the execution of the state machine.
5. *OutputCycleReset* — this operation allows the implementation to reset the output cycle state machine when unusual conditions occur.

OutputCycleInitialize initializes the *OutputCycle* state machine by moving it from its finished, to its initial state.

<i>OutputCycleInitialize</i>
ΔIO
$output_cycle_state = OutputCycleFinishedState$
$output_cycle_state' = OutputCycleInitialState$

StartSendOutputChannel starts the shipping of output by moving the machine into the *OutputCycleSendState* state and pushing an empty *CycleOfOutputs* onto the end of the *outputs* sequence.

<i>StartSendOutputChannel</i>
ΔIO
$output_cycle_state = OutputCycleInitialState$
$output_cycle_state' = OutputCycleSendState$
$outputs' = outputs \hat{\ } (\emptyset)$

If an *OutputLink* of the top level state does not have a corresponding *OutputStructure* in the last *CycleOfOutput*, then *SendOutputChannelOutputStructureObject* creates a new *OutputStructure* with the object of the link and adds it to the output.

<p style="text-align: center;"><u>SendOutputChannelOutputStructureObject</u></p> <p>ΔIO</p> <hr/> <p>$output_cycle_state = OutputCycleSendState = output_cycle_state'$</p> <p>$\exists i : OutputAttribute \bullet$</p> <p style="padding-left: 20px;">$\exists w : WM \mid w.id = TopState \wedge w.attribute = i \wedge$ $w.context_acceptable_preference = No \wedge$ $OutputMapping(i) \notin dom(last(outputs)) \bullet$</p> <p style="padding-left: 20px;">$\exists OS : OutputStructure \mid$ $OS.object = w.value \wedge OS.wmes = \emptyset \bullet$ $(last(outputs') = last(outputs) \oplus$ $\{ OutputMapping(i) \mapsto OS \} \wedge$ $front(outputs') = front(outputs))$</p>
--

SendOutputChannelOutputStructureWME finds any working memory elements of the transitive closure of an *OutputLink* that have not been added to their *OutputStructure* and adds them to the structure.

<p style="text-align: center;"><u>SendOutputChannelOutputStructureWME</u></p> <p>ΔIO</p> <hr/> <p>$output_cycle_state = OutputCycleSendState = output_cycle_state'$</p> <p>$\exists i : OutputAttribute \bullet$</p> <p style="padding-left: 20px;">$\exists w : WM \mid w.id = TopState \wedge w.attribute = i \wedge$ $w.context_acceptable_preference = No \bullet$</p> <p style="padding-left: 20px;">$\exists OS, OS' : OutputStructure \mid$ $OS.object = w.value \wedge OS'.object = w.value \bullet$ $\exists w_2 : OfId_{WME}(TCI_{WME}(WM)({w.value}), WM) \mid$ $w_2 \notin OS.wmes \wedge OS'.wmes = OS.wmes \cup \{w_2\} \bullet$ $(last(outputs') = last(outputs) \oplus$ $\{ OutputMapping(i) \mapsto OS' \} \wedge$ $front(outputs') = front(outputs))$</p>
--

SendOutputChannel composes the operations to send out *OutputStructures*.

$$SendOutputChannel \cong SendOutputChannelOutputStructureObject \vee SendOutputChannelOutputStructureWME$$

When there is no more output to send, *OutputCycleFinish* finishes the output cycle.

<i>OutputCycleFinish</i> ΔIO
<i>output_cycle_state</i> = <i>OutputCycleSendState</i> \neg pre <i>SendOutputChannel</i> <i>output_cycle_state'</i> = <i>OutputCycleFinishedState</i>

OutputCycleStep sequences all of the operations of the *OutputCycle*.

$$OutputCycleStep \cong StartSendOutputChannel \vee SendOutputChannel \vee OutputCycleFinish$$

OutputCycleReset allows the implementation to reset the *OutputCycle* state machine from any state under exceptional conditions.

<i>OutputCycleReset</i> ΔIO
<i>output_cycle_state'</i> = <i>OutputCycleInitialState</i>

Chapter 7

Decide

The decision procedure reads preference memory to decide what should be in working memory, including the context impasse stack and attribute impasses. Decide's specification is composed of a state schema and four state machines: impasser, preference semantics, working memory phase and quiescence phase. Soar's top level state machine steps decide through the working memory phase (WMPHase) or the quiescence phase (QPhase). The working memory phase and the quiescence phase step preference semantics to determine the values of slots and the impasser to add, modify and remove goals and impasses.

Z's bottom-up approach to specification requires that we first specify the state schema of a sequential system, and then specify its transitions. Consequently we present the state schemas for decide's four state machines first, then we present decide's total state schema, and then we present the state transitions for each of the component state machines.

This chapter is organized into ten sections.

- Section 7.1 Impasses — defines impasses, Soar's goal representation.
- Section 7.2 Impasser State — defines the state of the impasser.
- Section 7.3 Preference Semantics State — defines the state of preference semantics.
- Section 7.4 Working Memory Phase State — defines the state of working memory phase.
- Section 7.5 Quiescence Phase State — defines the state of quiescence phase.
- Section 7.6 Decide's Total State — defines the total state of decide.
- Section 7.7 Impasser — defines the state machine that implements impassing.

- Section 7.8 Preference Semantics — defines the state machine that implements preference semantics calculations.
- Section 7.9 Working Memory Phase — defines the state machine that implements the working memory phase.
- Section 7.10 Quiescence Phase — defines the state machine that implements quiescence phase.

7.1 Impasses

This section describes the basic concepts of an impasse, and provides some operations to help access and modify the working memory elements of impasses.

The context impasse stack is represented as elements in working memory, and so are its attribute impasses. Each element of a context impasse in the stack holds the context's identifier in its id component, and its attribute is an element of *ImpasseAttribute*, defined in Section 5.5.9.

$$\text{ImpassesWMEs} : \text{Symbol} \times \mathbb{P} \text{Identifier} \times \mathbb{P} \text{WME} \rightarrow \mathbb{P} \text{WME}$$

$$\forall g_or_i : \text{Symbol}; GM : \mathbb{P} \text{Identifier}; WM : \mathbb{P} \text{WME} \bullet$$

$$\text{ImpassesWMEs}(g_or_i, GM, WM) =$$

$$\{w : WM \mid g_or_i \in GM \wedge w.id = g_or_i \wedge \\ w.attribute \in \text{ImpasseAttribute} \wedge \\ w.context_acceptable_preference = \text{No}\}$$

There are eight attributes that are used to create working memory elements defining a context impasse.

1. 'PROBLEM-SPACE' — the optional problem space of the subgoal.
2. 'STATE' — the optional state of the subgoal.
3. 'OPERATOR' — the optional operator of the subgoal.
4. 'OBJECT' — the parent context impasse for a context impasse, or the identifier of the impassed slot for an attribute impasse.
5. 'ATTRIBUTE' — the attribute of the impassed slot.
6. 'IMPASSE' — the type of impasse (see 7.2).
7. 'ITEM' — the candidates that were available for decision for 'TIE', 'CONFLICT' or 'CONSTRAINT-FAILURE' impasse.
8. 'CHOICES' — if there are no 'ITEM' augmentations, an impasse has a 'CHOICES' augmentation of value 'NONE', otherwise its value is 'MULTIPLE'.

Decide uses only the 'OBJECT', 'ATTRIBUTE', 'IMPASSE', 'ITEM' and 'CHOICES' attributes for attribute impasses. The values from the ('OBJECT', 'ATTRIBUTE') working memory element pair form the slot for which decide had incomplete or inconsistent information. The 'IMPASSE' element holds as a value the type of impasse. The 'CHOICES' describes if there were 'MULTIPLE' or 'NONE' value choices available to the decision procedure for the impassed slot. Each choice is represented in working memory using an augmentation of attribute 'ITEM'.

The recursive function *GoalDescendants* follows the chain of 'OBJECT' attributes backwards through working memory to construct the set of identifiers of the descendant goals of a given goal.

$$\textit{GoalDescendants} : \textit{Identifier} \times \mathbb{P} \textit{Identifier} \times \mathbb{P} \textit{WME} \rightarrow \mathbb{P} \textit{Identifier}$$

$$\begin{aligned} &\forall g : \textit{Identifier}; GM : \mathbb{P} \textit{Identifier}; WM : \mathbb{P} \textit{WME} \bullet \\ &\quad (\exists_1 c : GM \bullet \textit{makeWME}(c, \textit{OBJECT}, g, No) \in WM) \Rightarrow \\ &\quad (\exists_1 c : GM \mid \textit{makeWME}(c, \textit{OBJECT}, g, No) \in WM \bullet \\ &\quad \quad \textit{GoalDescendants}(g, GM, WM) = \\ &\quad \quad \{c\} \cup \textit{GoalDescendants}(c, GM, WM)) \wedge \\ &\quad - (\exists_1 c : GM \bullet \textit{makeWME}(c, \textit{OBJECT}, g, No) \in WM) \Rightarrow \\ &\quad \quad \textit{GoalDescendants}(g, GM, WM) = \emptyset \end{aligned}$$

7.2 Impasser State

The impasser creates, modifies and deletes both context and attribute impasses. Its state machine requires only three states (Figure 7.1):

1. an initial state which is used to optionally remove old impasses
2. a state from which new impasses are created or old impasses have their items changed and
3. a final state.

$$\begin{aligned} \textit{ImpasserState} ::= & \textit{ImpasserInitialState} \\ & \mid \textit{ImpasserCreateOrChangeState} \\ & \mid \textit{ImpasserFinishedState} \end{aligned}$$

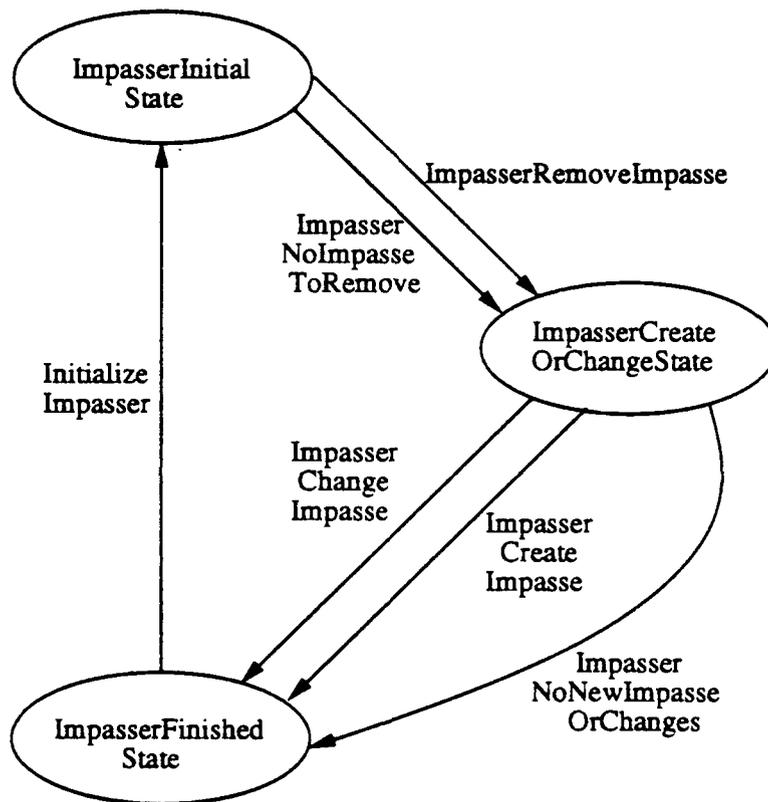


Figure 7.1: Impasser State Machine

The impasser is started with a slot that may need an impasse created, changed or removed, and the type of the impasse that is required.

<p><i>ImpasseType</i> : P Constant</p> <hr/> <p><i>ImpasseType</i> = { 'CONSTRAINT-FAILURE', 'CONFLICT', 'TIE', 'NO-CHANGE', 'NONE' }</p>

There are four types of impasses and one symbol used to specify that the slot does not require an impasse.

1. 'CONSTRAINT-FAILURE' - created when decide has more than one required candidate for a slot, or has a required candidate that is also prohibited.
2. 'CONFLICT' - created when decide has inconsistent better, worse, best or worst preferences for a slot.
3. 'TIE' - created when decide has more than one acceptable, non-parallel, non-indifferent candidate for a slot.
4. 'NO-CHANGE' - created when quiescence is reached and the decision procedure has no change that it can make to the context stack.
5. 'NONE' - signifies that preference semantics has determined that it has no candidate, or a unique candidate for a slot.

The flag, *context_changed*, of type *ContextChanged* is used to record if the context has been changed during quiescence phase or the impasser state machine. This allows the quiescence phase to know when the context has been changed so that it can stop looping over the slots of the context.

ContextChanged ::= *Changed* | *Unchanged*

The *ImpasserStateSchema* groups the slot to be impassed, the type of the impasse, the items under consideration from preference semantics and the state counter for the impasse state machine

<p><i>ImpasserStateSchema</i></p> <hr/> <p><i>impasser_state</i> : <i>ImpasserState</i> <i>slot</i> : <i>Slot</i> <i>impasse</i> : <i>ImpasseType</i> <i>items</i> : P <i>Symbol</i> <i>context_changed</i> : <i>ContextChanged</i></p>

7.3 Preference Semantics State

Preference semantics reads the contents of preference memory to decide new values for working memory slots. The preference semantics state machine requires eleven states (Figure 7.3): nine to perform the filter-like operations to decide slots, an initial and a final state.

1. **PSInitialState** — preference semantics starts out in this state.
2. **PSRequireState** — this state checks for required candidates in preference memory and also constraint failure impasses involving require preferences.
3. **PSAcceptableState** — this state collects the candidates that have acceptable preferences in memory.
4. **PSProhibitState** — this state filters out the candidates that have been prohibited.
5. **PSRejectState** — this state filters out the candidates that have rejects in PM.
6. **PSBetterWorseState** — this state checks for better worse conflicts, and otherwise filters out candidates based on better/worse preferences.
7. **PSBestState** — this state filters out candidates using the best preferences.
8. **PSWorstState** — this state filters out candidates using the worst preferences.
9. **PSIndifferentState** — If all of the remaining candidates are mutually indifferent, this state returns them as indifferent. If they are not, and this is a context impasse slot, a tie impasse is returned. Non-context slots's items are tested to see if they are mutually parallel.
10. **PSParallelState** — If all of the remaining candidates are mutually parallel, they're returned as mutually parallel. Otherwise, a tie impasse is returned.
11. **PSFinishedState** — preference semantics enters this state when it has reached a conclusion for its slot.

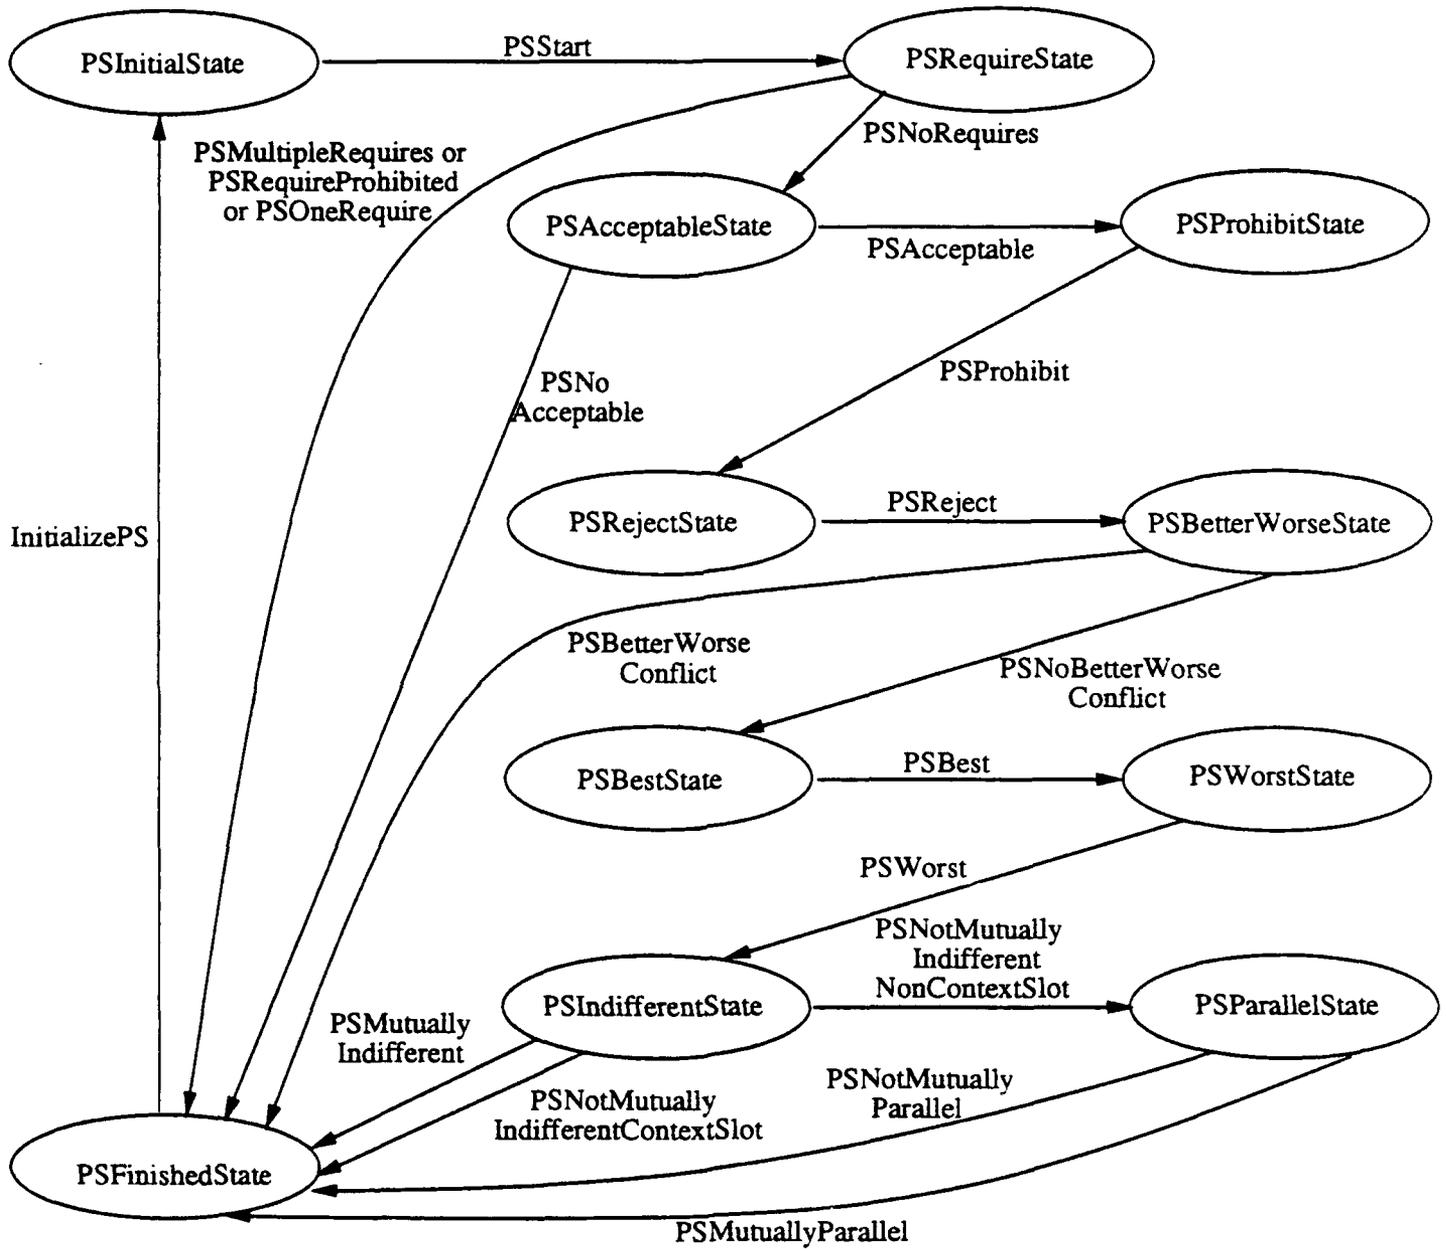


Figure 7.2: Preference Semantics State Machine

```

PreferenceSemanticsState ::= PSInitialState
                           PSRequireState
                           PSAcceptableState
                           PSProhibitState
                           PSRejectState
                           PSBetterWorseState
                           PSBestState
                           PSWorstState
                           PSIndifferentState
                           PSParallelState
                           PSFinishedState

```

Preference semantics requires an input slot, and produces an impasse, a possibly empty set of conflicting, indifferent or parallel items and a flag, *number_of_winners* of type *NumberOfWinners*. The *number_of_winners* flag is set whenever there is a winning value, e.g., there is not an impasse. If a single winner is chosen or there are a set of mutually indifferent winners, *number_of_winners* is *One*. If all of the winners are mutually parallel, then *number_of_winners* is *All*. This allows *WMPPhase* to distinguish a set of mutually parallel candidates that should all be installed in working memory from a set of mutually indifferent candidates that should have only one value installed or maintained.

```

NumberOfWinners ::= All | One

```

The *PreferenceSemanticsStateSchema* schema groups preference semantics's inputs, outputs and state counter.

```

PreferenceSemanticsStateSchema
ps_state : PreferenceSemanticsState
slot : Slot
items : P Symbol
impasse : Impasse Type
number_of_winners : NumberOfWinners

```

7.4 Working Memory Phase State

The *WMPPhase* state machine is implemented in six states (Figure 7.3).

1. *WMPPhaseInitialState* — the state machine starts out in this state.
2. *WMPPhasePickSlotState* — this state picks any non-context slot with changed preferences for decision.

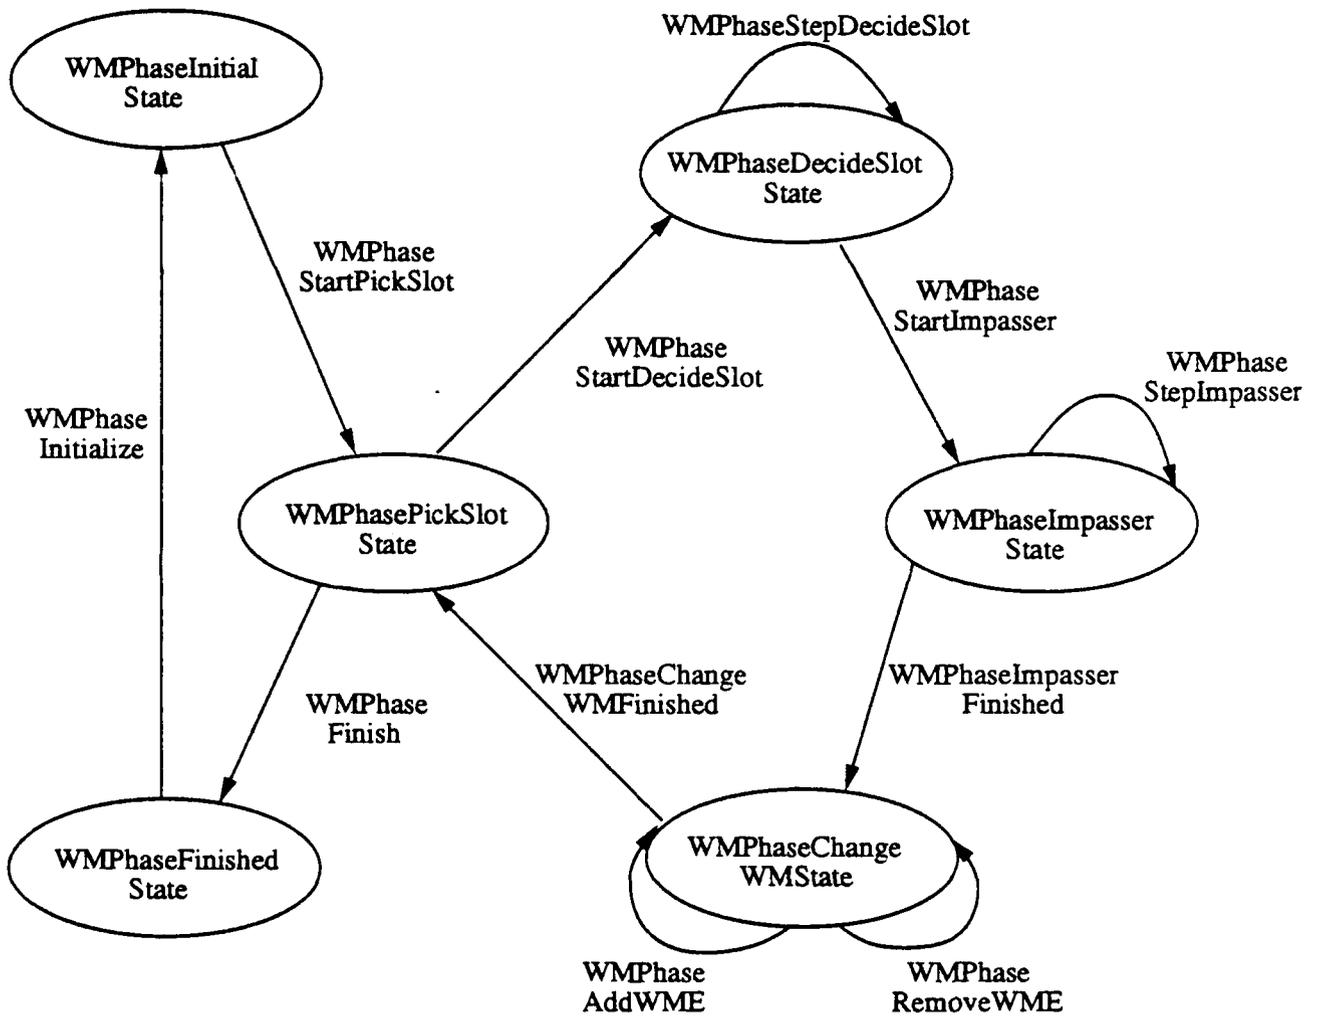


Figure 7.3: Working Memory Phase State Machine

3. *WMPhaseDecideSlotState* — this state steps the preference semantics state machine until it has arrived at a decision.
4. *WMPhaseImpasseState* — if the decision is an impasse, this state steps the impasse machine until it has added, changed or removed the impasse.
5. *WMPhaseChangeWMState* — this state adds or removes the WMEs for the slot.
6. *WMPhaseFinishedState* — the state machine finishes in this state.

```

WMPhaseState ::= WMPhaseInitialState
                WMPhasePickSlotState
                WMPhaseDecideSlotState
                WMPhaseImpasserState
                WMPhaseChangeWMState
                WMPhaseFinishedState

```

The preference phase records the set of slots that have preferences added or retracted in the *ChangedSlots* component of the *WMPhaseStateSchema*. The *WMPhase* destructively iterates the component *slot* over the changed slots, runs the decision procedure on the *slot* and receives back information in the *items*, *impasse* and the *number_of_winners* components.

```

WMPhaseStateSchema
-----
wmpPhase_state : WMPhaseState
changed_slots : P Slot
slot : Slot
items : P Symbol
impasse : ImpasseType
number_of_winners : NumberOfWinners

```

7.5 Quiescence Phase State

Quiescence Phase's state machine requires six states (Figure 7.4).

1. *QPhaseInitialState* — the state machine starts from this state.
2. *QPhasePickSlotState* — chooses the context it passes and which one of its slots to decide next.
3. *QPhaseDecideSlotState* — this state steps the preference semantics state machine given this slot as an argument.

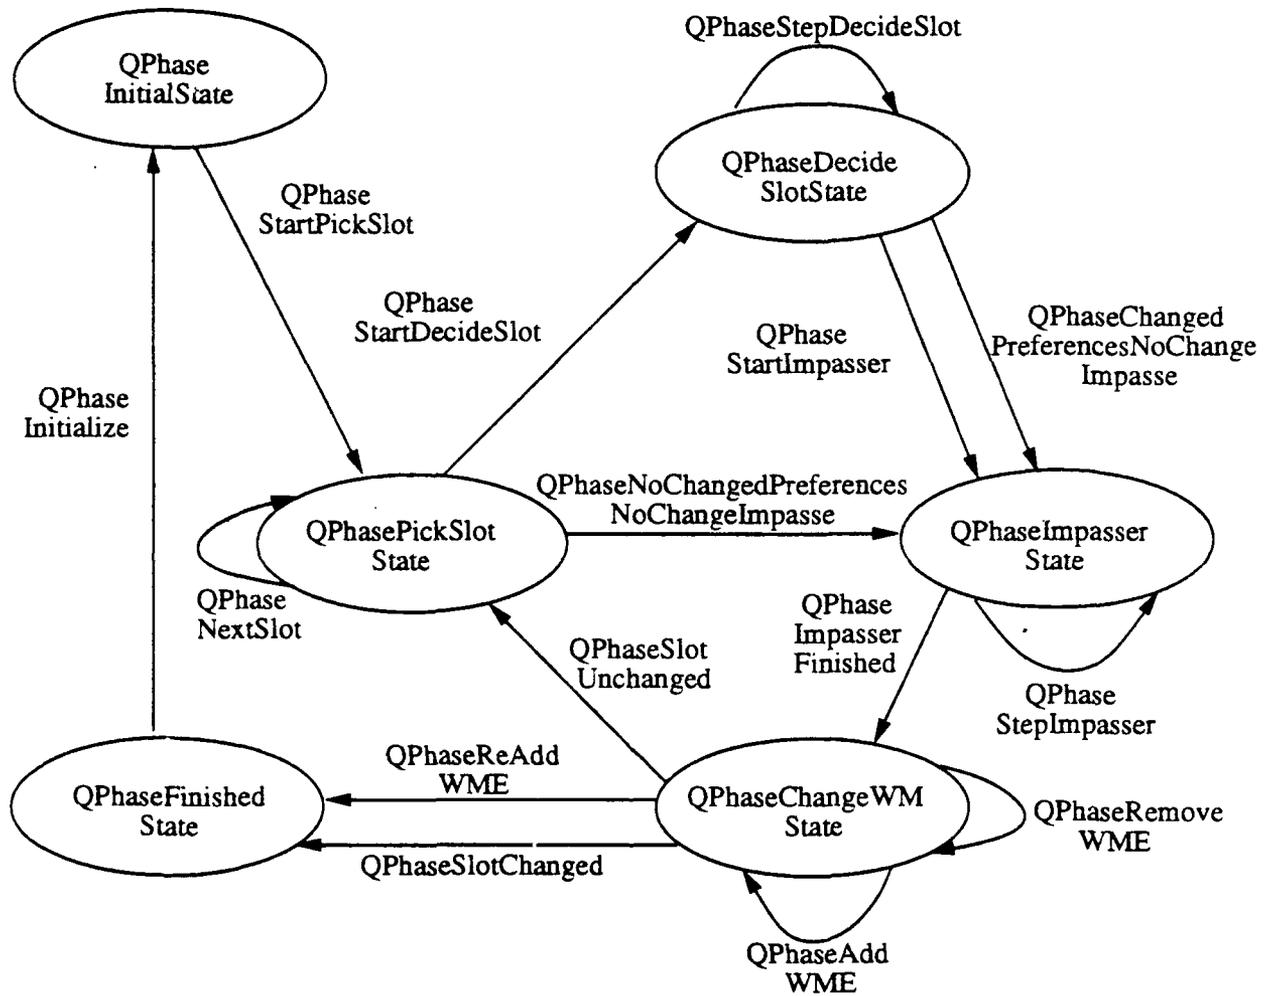


Figure 7.4: Quiescence Phase State Machine

4. *QPhaseImpasserState* — this state steps the impasser given the result of preference semantics.
5. *QPhaseChangeWMState* — this state changes the values in working memory for the slot.
6. *QPhaseFinishedState* — the final state of the machine.

```

QPhaseState ::= QPhaseInitialState
                | QPhasePickSlotState
                | QPhaseDecideSlotState
                | QPhaseImpasserState
                | QPhaseChangeWMState
                | QPhaseFinishedState

```

The set *ContextSlotAttribute* holds the attributes for context slots.

```

ContextSlotAttribute : P ImpasseAttribute
-----
ContextSlotAttribute =
  { `PROBLEM-SPACE`, `STATE`, `OPERATOR` }

```

The *QPhase* state machine uses six pieces of state:

1. a state counter, *qphase_state*
2. the slot of the context under consideration
3. a flag that marks if the context has been changed
4. the set of slots that preference phase has changed preferences for
5. the set of items that the preference semantics returned and
6. the number of winners amongst those items.

```

QPhaseStateSchema -----
qphase_state : QPhaseState
slot : Slot
context_changed : ContextChanged
changed_slots : P Slot
items : P Symbol
number_of_winners : NumberOfWinners

```

7.6 Decide's Total State

The decide state schema includes the state schemas for its state machines, defines the impasse memory, the goal memory and two pointers into it, and shares the temporary memory, *its* component memories and the identifier table with the other major modules of Soar.

The goal memory, GM, and impasse memory, IM, are defined as sets of identifiers. All information about the type of context impasse or attribute impasse is represented by working memory elements. The goal memory is also constrained to be contained within the transitive closure of the bottom goal. Impasse memory is constrained somewhat differently: all of the objects of the attribute impasses must be connected to the bottom goal.

Temporary memory, TM, is defined and constrained by three recursive axioms.

1. Temporary memory is the union of the temporary memory elements created from all of working memory and production memory, using the free type constructors for temporary memory elements.
2. Temporary memory is defined to be transitively closed from the bottom goal and the impasses. This means that all elements are on a path that starts with the identifier of the bottom goal or an impasse identifier.
3. Working memory, which may be viewed as a part of temporary memory, contains a working memory element of *context_acceptable_preference* Yes if and only if there is a corresponding acceptable preference for that context slot. The mapping *ContextImpassePreferences* is used to construct the set of working memory elements that correspond to the context acceptable preferences.

$$\text{ContextImpassePreferences} : \mathbb{P} \text{ Identifier} \times \mathbb{P} \text{ Preference} \rightarrow \mathbb{P} \text{ WME}$$

$$\forall GM : \mathbb{P} \text{ Identifier}; PM : \mathbb{P} \text{ Preference} \bullet$$

$$\text{ContextImpassePreferences}(GM, PM) =$$

$$\{p : \text{UnaryP} \sim \{PM \cap \text{ran UnaryP}\} \mid$$

$$p.id \in GM \wedge (p.preference = '+' \vee p.preference = '!') \wedge$$

$$p.attribute \in$$

$$\{\text{' PROBLEM-SPACE' }, \text{' STATE' }, \text{' OPERATOR' }\}$$

$$\bullet \text{makeWME}(p.id, p.attribute, p.value, \text{Yes})\}$$

Decide

IM : P Identifier
GM : P Identifier
bottom_goal, top_goal : Identifier
IdentifierTable : P Identifier
TM : P TME
WM : P WME
PM : P Preference
PreferenceSemanticsStateSchema
ImpasserStateSchema
WMPhaseStateSchema
QPhaseStateSchema

$TM = PreferenceTME(PM) \cup WMETME(WM)$
 $TM = Offd_{TME}((TCI_{TME}(TM))(\{bottom_goal\} \cup IM), TM)$
 $ContextImpassePreferences(GM, PM) =$
 $\{w : WM \mid w.context_acceptable_preference = Yes\}$
 $IM = \{i : IM \mid$
 $\exists o : Symbol \bullet$
 $WMETME(makeWME(i, 'OBJECT', o, No)) \in$
 $Offd_{TME}((TCI_{TME}(TM))(\{bottom_goal\} \cup IM), TM)\}$
 $GM = ((TCI_{TME}(TM))(\{bottom_goal\})) \cap GM$

Decide starts off with its four state machines in their finished state, and empty impasse and goal memories.

InitDecide

Decide

ps_state = PSFinishedState
impasser_state = ImpasserFinishedState
wmphase_state = WMPhaseFinishedState
qphase_state = QPhaseFinishedState
IM = \emptyset
GM = {}

The top level state machine initializes decide with the *DecideCreateFirstGoal* operation. It sets goal memory to a singleton, adds the object working memory element for this goal and its type working memory element, and sets the top and bottom goal pointers to the single goal.

<p><i>DecideCreateFirstGoal</i></p> <hr/> <p>Δ<i>Decide</i></p> <hr/> <p>$\exists g : Identifier \setminus IdentifierTable \bullet (GM = \{g\} \wedge$ $top_goal = g \wedge$ $bottom_goal = g \wedge$ $WM' = \{makeWME(g, 'OBJECT', 'NIL', No),$ $makeWME(g, 'TYPE', 'GOAL', No)\}$</p> <p><i>ps_state</i> = <i>PSFinishedState</i></p> <p><i>impasser_state</i> = <i>ImpasserFinishedState</i></p> <p><i>wmphase_state</i> = <i>WMPhaseFinishedState</i></p> <p><i>qphase_state</i> = <i>QPhaseFinishedState</i></p> <p><i>IM</i> = \emptyset</p>
--

7.7 Impasser Transitions

The impasser has six state transitions between its three states (Figure 7.1).

1. *ImpasserInitialize* — used to initialize the state machine.
2. *ImpasserNoImpasseToRemove* — If there is no need to remove an impasse, this transits the machine from the initial state to the *CreateOrChange* state.
3. *ImpasserRemoveImpasse* — If there is an old impasse in memory that is for the wrong type of impasse or the wrong slot, this transition removes it.
4. *ImpasserCreateImpasse* — If preference semantics has determined a need for an impasse that is not yet in memory, this creates it.
5. *ImpasserChangeImpasse* — If preference semantics has found a changed set of items for an existing 'TIE', 'CONFLICT' or 'CONSTRAINT-FAILURE' impasse, this transition changes the items set.
6. *ImpasserNoNewImpasseOrChanges* — If there is not a new impasse and no changes to an item set, this finishes the state machine's execution.

ImpasserInitialize readies the impasser to run by moving it from the final to the initial state.

<i>ImpasserInitialize</i> Δ Decide <hr/> <i>context_changed'</i> = <i>Unchanged</i> <hr/> <i>impasser_state'</i> = <i>ImpasserInitialState</i>
--

The impasser first checks for an old impasse existing in memory that must be removed. An impasse must be removed if preference semantics has decided that the slot needs a different type of impasse, a higher slot in the context has an impasse, or the slot should have no impasse at all. *ImpasserRemoveImpasse* is designed to remove both context and non-context impasses, so it searches for an impasse identifier in IM or GM. If the obsolete impasse is a non-context impasse, it deletes it from IM. If the obsolete impasse is a context impasse, it deletes it and its descendants from GM, and updates the pointer to the bottom of the context stack.

<i>ImpasserRemoveImpasse</i> Δ Decide <hr/> <i>impasser_state</i> = <i>ImpasserInitialState</i> $\exists i : IM \cup GM \mid$ $make\ WME(i, \text{'OBJECT'}, first(slot), No) \in WM \bullet$ $\neg ((make\ WME(i, \text{'ATTRIBUTE'}, second(slot), No) \in WM) \wedge$ $(make\ WME(i, \text{'IMPASSE'}, impasse.No) \in WM))$ \Rightarrow $((i \in IM \Rightarrow$ $IM' = IM \setminus \{i\} \wedge$ $WM' = WM \setminus ImpassesWMEs(i, IM, WM)) \wedge$ $(i \in GM \Rightarrow$ $(context_changed' = Changed \wedge$ $(GM' = GM \setminus (\{i\} \cup GoalDescendants(i, GM, WM))) \wedge$ $(bottom_goal' = first(slot))) \wedge$ $(WM' = WM \setminus ImpassesWMEs(i, GM, WM)$ $\setminus (\cup \{g : GoalDescendants(i, GM, WM) \bullet$ $ImpassesWMEs(g, GM, WM)\})))$ <hr/> <i>impasser_state'</i> = <i>ImpasserCreateOrChangeState</i>
--

If the impasse in memory is the desired type of impasse or there is no impasse in memory then *ImpasserNoImpasseToRemove* moves the impasser into the create or change state.

ImpasserNoImpasseToRemove

Δ Decide

impasser_state = *ImpasserInitialState*

$\exists i : IM \cup GM \bullet$

(*makeWME*(*i*, 'OBJECT', *first(slot)*, *No*) \in *WM* \wedge

makeWME(*i*, 'ATTRIBUTE', *second(slot)*, *No*) \in *WM* \wedge

makeWME(*i*, 'IMPASSE', *impasse*, *No*) \in *WM*) \vee

$\neg (\exists i : IM \cup GM \bullet$

(*makeWME*(*i*, 'OBJECT', *first(slot)*, *No*) \in *WM* \wedge

makeWME(*i*, 'ATTRIBUTE', *second(slot)*, *No*) \in *WM*))

impasser_state' = *ImpasserCreateOrChangeState*

If the desired impasse is not in memory, *ImpasserCreateImpasse* selects an identifier that is not in the *IdentifierTable*, and consequently not in use by the system, and adds the identifier to IM or GM and the basic three elements required for the impasse. If the impasse is a 'CONSTRAINT-FAILURE' or a 'NO-CHANGE' then an element of attribute 'CHOICES' value 'NONE' is also added. If the impasse is a 'CONFLICT' or a 'TIE' then an element of attribute 'CHOICES' value 'MULTIPLE' is added instead. An element of attribute 'ITEM' is added for each candidate returned by preference semantics in the items set. Chunking's *TraceItems* operation is referenced to give the new item elements numbers and setup their working memory traces to point to the preferences that caused them to be considered for the slot.

ImpasserCreateImpasse

$\Delta Decide$

impasser_state = *ImpasserCreateOrChangeState*

$\neg (\exists i : IM \cup GM \bullet$
 $\quad make WME(i, 'OBJECT', first(slot), No) \in WM \wedge$
 $\quad make WME(i, 'ATTRIBUTE', second(slot), No) \in WM)$

$\exists i : Identifier \mid i \notin IdentifierTable \bullet$
 $(first(slot) \in GM \wedge$
 $\quad second(slot) \in ContextSlotAttribute \cup \{ 'GOAL' \}) \Rightarrow$
 $\quad (GM' = GM \cup \{i\} \wedge$
 $\quad\quad context_changed' = Changed \wedge$
 $\quad\quad bottom_goal' = i) \wedge$
 $\neg (first(slot) \in GM \wedge$
 $\quad second(slot) \in ContextSlotAttribute \cup \{ 'GOAL' \}) \Rightarrow$
 $\quad IM' = IM \cup \{i\} \wedge$

$\exists W : P WME \mid W = \{v : items \bullet make WME(i, 'ITEM', v, No)\}$
 $WM' = WM \cup$
 $\quad \{make WME(i, 'OBJECT', first(slot), No),$
 $\quad\quad make WME(i, 'ATTRIBUTE', second(slot), No),$
 $\quad\quad make WME(i, 'IMPASSE', impasse, No)\} \cup$
 $\quad \{w : WME \mid w = make WME(i, 'CHOICES', 'NONE', No) \wedge$
 $\quad\quad\quad impasse \in \{ 'CONSTRAINT-FAILURE', 'NO-CHANGE' \}\} \cup$
 $\quad \{w : WME \mid w = make WME(i, 'CHOICES', 'MULTIPLE', No) \wedge$
 $\quad\quad\quad impasse \in \{ 'CONFLICT', 'TIE' \}\} \cup W \wedge$
 $TraceItems$

impasser_state' = *ImpasserFinishedState*

If the correct type of impasse for the slot exists in memory, but its set of items is now different, *ChangeImpasse* removes the elements for the items that are no longer involved in the impasse, and adds in elements for the new items. This does not change the main structure of the context stack, only its items, so the *context_changed* flag is not set and sub-slots and sub-goals are not flushed. Having updated working memory to match preference semantic's required impasse, the *ChangeImpasse* steps to the finished *impasser* state.

ImpasserChangeImpasse

Δ Decide

impasser_state = *ImpasserCreateOrChangeState*

$\exists i : IM \cup GM \mid$
 $makeWME(i, \text{'OBJECT'}, first(slot), No) \in WM \wedge$
 $makeWME(i, \text{'ATTRIBUTE'}, second(slot), No) \in WM \wedge$
 $makeWME(i, \text{'IMPASSE'}, impasse, No) \in WM \wedge$
 $items \neq \{v : Symbol \mid makeWME(i, \text{'ITEM'}, v, No) \in WM\} \bullet$

$\exists W : P WME \mid$
 $W = \{w : WM \mid$
 $\exists v : Symbol \setminus items \bullet w = makeWME(i, \text{'ITEM'}, v, No)\}$
 $\cup \{v : items \mid makeWME(i, \text{'ITEM'}, v, No) \notin WM \bullet$
 $makeWME(i, \text{'ITEM'}, v, No)\} \bullet$
 $WM' = WM \setminus W \wedge TraceItems$

impasser_state' = *ImpasserFinishedState*

If there is no need for a new impasse, or the impasse in memory is exactly right, then *ImpasserNoNewImpasseOrChanges* completes the *impasser's* run.

ImpasserNoNewImpasseOrChanges

Δ Decide

impasser_state = *ImpasserCreateOrChangeState*

impasse = 'NONE' \vee
 $(\exists i : IM \cup GM \bullet$
 $makeWME(i, \text{'OBJECT'}, first(slot), No) \in WM \wedge$
 $makeWME(i, \text{'ATTRIBUTE'}, second(slot), No) \in WM \wedge$
 $makeWME(i, \text{'IMPASSE'}, impasse, No) \in WM \wedge$
 $items = \{v : Symbol \mid makeWME(i, \text{'ITEM'}, v, No) \in WM\})$

impasser_state' = *ImpasserFinishedState*

After initialization *ImpasserStep* will drive the impasser state machine through all of its states to completion.

$$\begin{aligned} \textit{ImpasserStep} \hat{=} & \textit{ImpasserNoImpasseToRemove} \vee \textit{ImpasserRemoveImpasse} \vee \\ & \textit{ImpasserCreateImpasse} \vee \textit{ImpasserChangeImpasse} \vee \\ & \textit{ImpasserNoNewImpasseOrChanges} \end{aligned}$$

7.8 Preference Semantics Transitions

Preference Semantics is defined in seventeen transitions (Figure 7.3): one to initialize, fifteen for processing the filter-like actions and one that is a concession to the obvious implementation to allow a fast exit from the state machine in a common case.

The *PSInitialize* initializes the state counter, empties the set of candidate items, and sets the impasse to 'NONE'.

$$\begin{array}{l} \textit{PSInitialize} \\ \hline \Delta \textit{Decide} \\ \hline \textit{ps_state}' = \textit{PSInitialState} \\ \textit{items}' = \emptyset \\ \textit{impasse}' = \text{'NONE'} \\ \hline \end{array}$$

StartPS moves the state machine into the state to test for require preferences.

$$\begin{array}{l} \textit{PSStart} \\ \hline \Delta \textit{Decide} \\ \hline \textit{ps_state} = \textit{PSInitialState} \\ \textit{slot}' = \textit{slot} \\ \textit{ps_state}' = \textit{PSRequireState} \\ \hline \end{array}$$

If there is exactly one require preference for the slot in memory and there is no prohibit preference for it in preference memory, then its value is the winner. The impasse is marked as 'NONE', the value is returned in the singleton items set, and preference semantics is finished.

<p><i>PSOneRequire</i></p> <hr/> <p>Δ<i>Decide</i></p> <hr/> <p><i>ps_state</i> = <i>PSRequireState</i></p> <p>$(\exists_1 v : \text{Symbol} \bullet$ $\text{makeUnaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), v, '!') \in PM) \wedge$ $\neg (\exists v : \text{Symbol} \bullet$ $\text{makeUnaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), v, '!') \in PM \wedge$ $\text{makeUnaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), v, '-') \in PM)$</p> <p>$\exists_1 v : \text{Symbol} \mid$ $\text{makeUnaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), v, '!') \in PM \bullet$ $\text{items}' = \{v\}$</p> <p><i>impasse'</i> = 'NONE'</p> <p><i>number_of_winners'</i> = <i>One</i></p> <p><i>ps_state'</i> = <i>PSFinishedState</i></p>
--

If there is more than one required candidate, then a constraint failure impasse is recognized with all of the required items in the items set, and preference semantics is complete.

<p><i>PSMultipleRequires</i></p> <hr/> <p>Δ<i>Decide</i></p> <hr/> <p><i>ps_state</i> = <i>PSRequireState</i></p> <p>$\exists V : \mathcal{P} \text{Symbol} \mid$ $(V = \{v : \text{Symbol} \mid$ $\text{makeUnaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), v, '!') \in PM\} \wedge$ $\#V > 1) \bullet$ $\text{items}' = V$</p> <p><i>impasse'</i> = 'CONSTRAINT-FAILURE'</p> <p><i>ps_state'</i> = <i>PSFinishedState</i></p>

If there exists a required value that is also prohibited, a constraint impasse with the required/prohibited value is recognized and preference semantics terminates.

<i>PSRequireProhibited</i>
$\Delta Decide$
$ps_state = PSRequireState$
$\exists v : Symbol$
$makeUnaryPreference(first(slot), second(slot), v, '!') \in PM \wedge$
$makeUnaryPreference(first(slot), second(slot), v, '~') \in PM \bullet$
$items' = \{v\}$
$impasse' = \text{'CONSTRAINT-FAILURE'}$
$ps_state' = PSFinishedState$

If there are no require preferences, preference semantics continues on to consider acceptable preferences.

<i>PSNoRequires</i>
$\Delta Decide$
$ps_state = PSRequireState$
$\neg (\exists v : Symbol \bullet$
$makeUnaryPreference(first(slot), second(slot), v, '!') \in PM)$
$impasse' = \text{'NONE'}$
$items' = \emptyset$
$ps_state' = PSAcceptableState$

The require processing is stepped through by the *Require* transition. The *PSRequire* compound transition is defined to allow users a single place to hook all possible ways that preference semantics processes requires.

$$PSRequire \cong PSOneRequire \vee PSMultipleRequires \vee PSRequireProhibited \vee PSNoRequires$$

The *PSAcceptable* transition collects all candidates with acceptable preferences in preference memory, stores them in the items set and moves the machine to the state that checks for prohibit preferences.

<i>PSAcceptable</i>
Δ Decide
$ps_state = PSAcceptableState$
$\exists v : Symbol \bullet$ $makeUnaryPreference(first(slot), second(slot), v, '+') \in PM$
$slot' = slot$
$items' = \{v : Symbol \mid$ $makeUnaryPreference(first(slot), second(slot), v, '+') \in PM\}$
$impasse' = \text{'NONE'}$
$ps_state' = PSProhibitState$

The *PSNoAcceptable* transition causes preference semantics to exit with an empty items set when there are no acceptable preferences in memory for the slot. This transition is a compromise to the obvious implementation. Preference semantics would run exactly the same without it, but one of the common cases is re-deciding a slot that has had all of its acceptable preferences withdrawn and so every implementation would benefit from optimizing this case.

<i>PSNoAcceptable</i>
Δ Decide
$ps_state = PSAcceptableState$
$\neg (\exists v : Symbol \bullet$ $makeUnaryPreference(first(slot), second(slot), v, '+') \in PM)$
$slot' = slot$
$items' = \emptyset$
$impasse' = \text{'NONE'}$
$ps_state' = PSFinishedState$

The *PSProhibit* transition filters out of the current set of candidates those which have prohibit preferences, and moves the machine to a state that searches for reject preferences.

<p><i>PSProhibit</i></p> <hr/> <p>Δ<i>Decide</i></p> <hr/> <p><i>ps_state</i> = <i>PSProhibitState</i></p> <p><i>slot'</i> = <i>slot</i></p> <p><i>items'</i> = { <i>v</i> : <i>items</i> $\text{makeUnaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), v, \text{'-'}) \notin PM$ }</p> <p><i>impasse'</i> = 'NONE'</p> <p><i>ps_state'</i> = <i>PSRejectState</i></p>

The *PSReject* transition filters out of the set of items those that have reject preferences, and jumps to processing the better worse preferences.

<p><i>PSReject</i></p> <hr/> <p>Δ<i>Decide</i></p> <hr/> <p><i>ps_state</i> = <i>PSRejectState</i></p> <p><i>slot'</i> = <i>slot</i></p> <p><i>items'</i> = { <i>v</i> : <i>items</i> $\text{makeUnaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), v, \text{'-'}) \notin PM$ }</p> <p><i>impasse'</i> = 'NONE'</p> <p><i>ps_state'</i> = <i>PSBetterWorseState</i></p>
--

The function *Conflict* is used to extract the candidates that have conflicting better/worse preferences.

$$\text{Conflict} : \mathbb{P} \text{ Preference} \times \text{Slot} \times \mathbb{P} \text{ Symbol} \rightarrow \mathbb{P} \text{ Symbol}$$

$$\forall PM : \mathbb{P} \text{ Preference}; \text{slot} : \text{Slot}; V : \mathbb{P} \text{ Symbol} \bullet$$

$$\text{Conflict}(PM, \text{slot}, V) =$$

$$\{j : V \mid \exists k : V \mid j \neq k \bullet$$

$$(\text{makeBinaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), j, '>', k) \in PM \wedge$$

$$\text{makeBinaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), k, '>', j) \in PM) \vee$$

$$(\text{makeBinaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), j, '<', k) \in PM \wedge$$

$$\text{makeBinaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), k, '<', j) \in PM) \vee$$

$$(\text{makeBinaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), j, '>', k) \in PM \wedge$$

$$\text{makeBinaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), j, '<', k) \in PM)\}$$

If there are no conflicting better/worse preferences between the items, then *PSNoBetterWorseConflict* applies. It filters out of the items set the ones that have another item that is better, or are worse than another item. The state machine moves on to filter by best/worst preferences.

$$\text{PSNoBetterWorseConflict}$$

$$\Delta \text{Decide}$$

$$ps_state = \text{PSBetterWorseState}$$

$$slot' = slot$$

$$\text{Conflict}(PM, \text{slot}, \text{items}) = \emptyset$$

$$\text{items}' = \{v : \text{items} \mid$$

$$\neg (\exists w : \text{items} \bullet v \neq w \wedge$$

$$\text{makeBinaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), v, '<', w) \in PM \wedge$$

$$\text{makeBinaryPreference}(\text{first}(\text{slot}), \text{second}(\text{slot}), w, '>', v) \in PM)\}$$

$$\text{impasse}' = \text{'NONE'}$$

$$ps_state' = \text{PSBestState}$$

If there are better/worse conflicts, this transition sets the items augmentation to the a' of the items involved in the conflicts, declares the impasse to be a conflict and finishes preference semantics.

<p><i>PSBetterWorseConflict</i></p> <hr/> <p>ΔDecide</p> <hr/> <p>$ps_state = PSBetterWorseState$</p> <p>$slot' = slot$</p> <p>$Conflict(PM, slot, items) \neq \emptyset$</p> <p>$impasse = \text{'CONFLICT'}$</p> <p>$items' = Conflict(PM, slot, items)$</p> <p>$ps_state' = PSFinishedState$</p>

The processing of better/worse preferences occurs in either the *PSNoBetterWorseConflict* or the *PSBetterWorseConflict* transitions and is stepped by *BetterWorse*.

$$PSBetterWorse \hat{=} PSNoBetterWorseConflict \vee PSBetterWorseConflict$$

The *PSBest* transition filters the candidates for those that have a best preference in memory. If there are no best preferences in memory for any of the current candidates, the entire old set of candidates is passed on. After best candidate filtering, preference semantics uses the worst preferences as a filter.

<p><i>PSBest</i></p> <hr/> <p>ΔDecide</p> <hr/> <p>$ps_state = PSBestState$</p> <p>$slot' = slot$</p> <p>$\exists V : P \text{ Symbol} \mid$ $V = \{v : items \mid$ $\text{makeUnaryPreference}(\text{first}(slot), \text{second}(slot), v, '>') \in PM\} \bullet$ $(V \neq \emptyset \Rightarrow items' = V \wedge$ $V = \emptyset \Rightarrow items' = items)$</p> <p>$impasse' = \text{'NONE'}$</p> <p>$ps_state' = PSWorstState$</p>
--

The *PSWorst* filter removes the items that have worst preferences in memory. If they all have worst preferences then none are filtered. After worst processing, preference semantics studies the indifferent preferences.

$\overline{PSWorst}$ $\Delta Decide$ <hr/> $ps_state = PSWorstState$ $slot' = slot$ $\exists V : \mathbb{P} Symbol \mid$ $V = \{v : items \mid$ $makeUnaryPreference(first(slot), second(slot), v, '<') \notin PM\} \bullet$ $(V \neq \emptyset \Rightarrow items' = V \wedge$ $V = \emptyset \Rightarrow items' = items)$ $impasse' = 'NONE'$ $ps_state' = PSIndifferentState$

The predicate *AllMutuallyIndifferent* checks that all of the candidates in the set are mutually indifferent. A candidate is mutually indifferent with the other candidates if it is unarily indifferent, or if there is a binary indifferent preference between the two candidates.

$\overline{AllMutuallyIndifferent_} : \mathbb{P}(\mathbb{P} Preference \times Slot \times \mathbb{P} Symbol)$ $\forall PM : \mathbb{P} Preference; slot : Slot; V : \mathbb{P} Symbol \bullet$ $AllMutuallyIndifferent(PM, slot, V) \Leftrightarrow$ $(\forall j, k : V : j \neq k \bullet$ $(makeUnaryPreference(first(slot), second(slot), j, '=p') \in PM \vee$ $((makeBinaryPreference(first(slot), second(slot), j, '=p', k) \in PM) \vee$ $makeBinaryPreference(first(slot), second(slot), k, '=p', j) \in PM)))$
--

If the candidates are all mutually indifferent, preference semantics returns with all of the candidates in the items set, an impasse of 'NONE' and sets the number of winners to one. This allows WMPPhase and QPhase to tell that it has a set of indifferent candidates to choose one from.

<i>PSMutuallyIndifferent</i>
Δ Decide
<i>ps_state</i> = <i>PSIndifferentState</i>
<i>slot'</i> = <i>slot</i>
<i>impasse</i> = 'NONE'
<i>AllMutuallyIndifferent</i> (<i>PM</i> , <i>slot</i> , <i>items</i>)
<i>items'</i> = <i>items</i>
<i>ps_state'</i> = <i>PSFinishedState</i>
<i>number_of_winners'</i> = <i>One</i>

If the candidates are not mutually indifferent and the decision is not for a context slot, there is still a chance that all of the items are parallel. In this case, *PSNotMutuallyIndifferentNonContextSlot* passes the items on to the parallel checking state.

<i>PSNotMutuallyIndifferentNonContextSlot</i>
Δ Decide
<i>ps_state</i> = <i>PSIndifferentState</i>
<i>slot'</i> = <i>slot</i>
<i>impasse</i> = 'NONE'
- <i>AllMutuallyIndifferent</i> (<i>PM</i> , <i>slot</i> , <i>items</i>)
- (<i>first</i> (<i>slot</i>) \in <i>GM</i> \wedge <i>second</i> (<i>slot</i>) \in { 'PROBLEM-SPACE', 'STATE', 'OPERATOR' })
<i>items'</i> = <i>items</i>
<i>ps_state'</i> = <i>PSParallelState</i>

Soar6 does not support parallel candidates for context items, although an earlier version of Soar did for operators. So non-mutually indifferent items for context slots generate a tie, with all of the candidates as items, not just those that are not mutually indifferent.

<p><i>PSNotMutuallyIndifferentContextSlot</i></p> <p>Δ<i>Decide</i></p> <p><i>ps_state</i> = <i>PSIndifferentState</i></p> <p><i>slot'</i> = <i>slot</i></p> <p><i>impasse</i> = 'NONE'</p> <p>\neg <i>AllMutuallyIndifferent</i>(<i>PM</i>, <i>slot</i>, <i>items</i>)</p> <p><i>first</i>(<i>slot</i>) \in <i>GM</i> \wedge</p> <p><i>second</i>(<i>slot</i>) \in { 'PROBLEM-SPACE', 'STATE', 'OPERATOR' }</p> <p><i>items'</i> = <i>items</i></p> <p><i>ps_state'</i> = <i>PSFinishedState</i></p> <p><i>impasse'</i> = 'TIE'</p>
--

The complete indifference processing is handled by one of the three transitions above, and stepped and hooked using the *Indifferent* transition.

PSIndifferent $\hat{=}$ *PSMutuallyIndifferent* \vee
PSNotMutuallyIndifferentNonContextSlot \vee
PSNotMutuallyIndifferentContextSlot

The parallel state uses the predicate *AllMutuallyParallel* to check if the set of candidates are all mutually parallel. An item is mutually parallel to the set of other items if it has a unary parallel preference, or if it has either of the two possible relative parallels between each item.

<p><i>AllMutuallyParallel</i> : $\mathbb{P}(\mathbb{P} \text{ Preference} \times \text{Slot} \times \mathbb{P} \text{ Symbol})$</p> <p>$\forall PM : \mathbb{P} \text{ Preference}; \text{slot} : \text{Slot}; V : \mathbb{P} \text{ Symbol} \bullet$</p> <p><i>AllMutuallyParallel</i>(<i>PM</i>, <i>slot</i>, <i>V</i>) \Leftrightarrow</p> <p>($\forall j, k : V \mid j \neq k \bullet$</p> <p>(<i>makeUnaryPreference</i>(<i>first</i>(<i>slot</i>), <i>second</i>(<i>slot</i>), <i>j</i>, '&') \in <i>PM</i> \vee</p> <p><i>makeBinaryPreference</i>(<i>first</i>(<i>slot</i>), <i>second</i>(<i>slot</i>), <i>j</i>, '&', <i>k</i>) \in <i>PM</i> \vee</p> <p><i>makeBinaryPreference</i>(<i>first</i>(<i>slot</i>), <i>second</i>(<i>slot</i>), <i>k</i>, '&', <i>j</i>) \in <i>PM</i>))</p>
--

If all of the candidates are mutually parallel *PSMutuallyParallel* returns all of the items, sets impasse to 'NONE' and the number of winners to All.

PSMutuallyParallel

Δ Decide

ps_state = *PSParallelState*
slot' = *slot*
AllMutuallyParallel(*PM*, *slot*, *items*)
items' = *items*
impasse' = 'NONE'
number_of_winners' = All
ps_state' = *PSFinishedState*

PSNotMutuallyParallel returns a tie impasse if all of the candidates are not mutually parallel.

PSNotMutuallyParallel

Δ Decide

ps_state = *PSParallelState*
slot' = *slot*
 \neg *AllMutuallyParallel*(*PM*, *slot*, *items*)
items' = *items*
impasse' = 'TIE'
ps_state' = *PSFinishedState*

The *PSParallel* transition is defined to either be a *MutuallyParallel* transition or a *NotMutuallyParallel* transition.

$$PSParallel \hat{=} PSMutuallyParallel \wedge PSNotMutuallyParallel$$

The operation to step preference semantics executes the start transition, or one of the transitions associated with the processing for a specific preference.

$$PSStep \hat{=} PSStart \vee PSRequire \vee PSAcceptable \vee PSNoAcceptable \vee PSProhibit \vee PSReject \vee PSBetterWorse \vee PSBest \vee PSWorst \vee PSIndifferent \vee PSParallel$$

7.9 Working Memory Phase Transitions

Working Memory Phase is defined in 12 transitions (Figure 7.3).

1. *WMPPhaseInitialize* — prepares the WMPPhase state machine for execution.
2. *WMPPhaseReset* — resets the state of WMPPhase when its execution is terminated abnormally.
3. *WMPPhaseStartPickSlot* — starts the loop over the changed slots.
4. *WMPPhaseStartDecideSlot* — selects one of the changed slots to decide.
5. *WMPPhaseStepDecideSlot* — runs the preference semantics on the chosen slot.
6. *WMPPhaseStartImpasser* — starts the execution of the impasser on the slot.
7. *WMPPhaseStepImpasser* — steps the impasser state machine.
8. *WMPPhaseImpasserFinished* — when the impasser is finished, this moves the machine on to the *WMPPhaseChangeWMState* which adds and removes working memory elements for the slot.
9. *WMPPhaseRemoveWME* — this removes any working memory elements no longer supported by the slot's decision.
10. *WMPPhaseAddWME* — this adds in any working memory elements that are required by the slot's decision.
11. *WMPPhaseChangeWMFinished* — when all of the required changes to working memory have been accomplished, this returns control to selecting another changed slot.
12. *WMPPhaseFinish* — when all of the changed non-context slots have been decided, this stops the execution of WMPPhase.

WMPPhaseInitialize initializes the WMPPhase by resetting the state counter.

<i>WMPPhaseInitialize</i> _____ Δ <i>Decide</i>
<i>wmphase_state</i> = <i>WMPPhaseFinishedState</i>
<i>wmphase_state'</i> = <i>WMPPhaseInitialState</i>

WMPHaseReset resets the system if *WMPHase*'s operation is abnormally terminated from the external interface.

<i>WMPHaseReset</i> _____ Δ <i>Decide</i>
<i>wmphase_state'</i> = <i>WMPHaseInitialState</i>

WMPHaseStartPickSlot moves *WMPHase* into the state that will select a changed slot to decide.

<i>WMPHaseStartPickSlot</i> _____ Δ <i>Decide</i>
<i>wmphase_state</i> = <i>WMPHaseInitialState</i>
<i>wmphase_state'</i> = <i>WMPHasePickSlotState</i>

The recognition memory records all of the slots that had preference changes during the last preference phase in the set *changed_slots*. In the pick slot state *WMPHase* picks one of these slots to have its preferences decided, initializes the preference semantics and moves to the state that decides the slot.

<i>WMPHaseStartDecideSlot</i> _____ Δ <i>Decide</i>
<i>wmphase_state</i> = <i>WMPHasePickSlotState</i>
$\exists s : \text{changed_slots} \setminus (GM \times \text{ContextSlotAttribute}) \bullet$ $(\text{changed_slots}' = \text{changed_slots} \setminus \{s\} \wedge$ $\text{slot}' = s)$
<i>PSInitialize</i>
$\text{slot}' = \text{slot}$
<i>wmphase_state'</i> = <i>WMPHaseDecideSlotState</i>

The *WMPHaseStepDecideSlot* transition steps preference semantics until a decision has been reached.

<i>WMPHaseStepDecideSlot</i> $\Delta Decide$ <hr/> <i>wmphase_state = WMPHaseDecideSlotState = wmphase_state'</i> <i>PSStep</i> <i>slot' = slot</i>

When a decision has been reached preference semantic's precondition is no longer satisfied, so *WMPHaseStartImpasser* initializes the impasser and moves *WMPHase* to the *WMPHaseImpasserState* state to run the impasser.

<i>WMPHaseStartImpasser</i> $\Delta Decide$ <hr/> <i>wmphase_state = WMPHaseDecideSlotState</i> - pre <i>PSStep</i> <i>ImpasserInitialize</i> <i>slot' = slot</i> <hr/> <i>wmphase_state' = WMPHaseImpasserState</i>
--

The *WMPHaseStepImpasser* transition runs the impasser. The impasser looks at the decision that preference semantics has created, and changes working memory and impasse memory to match. If the decision requires an impasse, the impasser will create it. If the decision requires a change to the existing impasse, the impasser will change it. If the decision requires no impasse, the impasser will remove any existing impasse. If the decision requires no impasse and none exists in memory, the impasser will make no changes.

<i>WMPHaseStepImpasser</i> $\Delta Decide$ <hr/> <i>wmphase_state = WMPHaseImpasserState = wmphase_state'</i> <i>ImpasserStep</i> <i>slot' = slot</i>

The *WMPhaseImpasserFinished* transition recognizes when the impasser has finished by checking its precondition. When it has completed, it moves the *WMPhase* on to the state that changes the values in working memory to match the slot's decision.

<i>WMPhaseImpasserFinished</i>
Δ <i>Decide</i>
<i>wmphase_state</i> = <i>WMPhaseImpasserState</i>
- pre <i>ImpasserStep</i>
<i>slot'</i> = <i>slot</i>
<i>wmphase_state'</i> = <i>WMPhaseChangeWMState</i>

If there is no impasse, and there is not some element in working memory that the decision requires then the *WMPHaseAddWME* transition will add it. If the decision is for a set of all parallel items, then *WMPHaseAddWME* will add them all, one at a time. If the decision is for one of a set of indifferent candidates and one candidate is not already installed, then *WMPHaseAddWME* will pick one of the items non-deterministically and add it to memory. Chunking's *TraceWME* operation is referenced to extend the instance numbering and the working memory trace.

<p><i>WMPHaseAddWME</i></p> <p>Δ <i>Decide</i></p> <hr/> <p><i>wmphase_state</i> = <i>WMPHaseChangeWMState</i> = <i>wmphase_state'</i></p> <p><i>slot'</i> = <i>slot</i></p> <p><i>items'</i> = <i>items</i></p> <p><i>impasse</i> = 'NONE'</p> <p>$\exists v : \textit{items} !$</p> <p style="padding-left: 2em;">(<i>number_of_winners</i> = <i>All</i> \wedge</p> <p style="padding-left: 4em;">$\neg (\exists w : \textit{WM} \bullet$</p> <p style="padding-left: 6em;"><i>w.id</i> = <i>first(slot)</i> \wedge <i>w.attribute</i> = <i>second(slot)</i> \wedge</p> <p style="padding-left: 6em;"><i>w.value</i> = <i>v</i> \wedge <i>w.context_acceptable_preference</i> = <i>No</i>)</p> <p style="padding-left: 4em;">)\vee</p> <p style="padding-left: 2em;">(<i>number_of_winners</i> = <i>One</i> \wedge</p> <p style="padding-left: 4em;">$\neg (\exists w : \textit{WM} \bullet$</p> <p style="padding-left: 6em;"><i>w.id</i> = <i>first(slot)</i> \wedge <i>w.attribute</i> = <i>second(slot)</i> \wedge</p> <p style="padding-left: 6em;"><i>w.value</i> \in <i>items</i> \wedge <i>w.context_acceptable_preference</i> = <i>No</i>)</p> <p style="padding-left: 4em;">)\bullet</p> <p>$\exists w : \textit{WME} \mid w = \textit{makeWME}(\textit{first}(\textit{slot}), \textit{second}(\textit{slot}), v, \textit{No}) \wedge$</p> <p><i>WM'</i> = <i>WM</i> \cup {<i>w</i>} \wedge</p> <p><i>TraceWME</i>)</p>
--

If there is an element in working memory that does not match the decision, the *WMPHaseRemoveWME* transition removes it. If the decision resulted in an impasse, then *WMPHaseRemoveWME* removes all of the working memory elements for the slot. If the decision resulted in no impasse, and there is a value in memory that is not in the items set, then it is removed. If the decision is for a set of mutually indifferent candidates and more than one of them is in memory, then all but one are chosen non-deterministically to removed.

<i>WMPHaseRemoveWME</i>
Δ <i>Decide</i>
$wmphase_state = WMPHaseChangeWMState = wmphase_state'$
$slot' = slot$
$items' = items$
$\exists w : WM \mid$
$w.id = first(slot) \wedge w.attribute = second(slot) \wedge$
$w.context_acceptable_preference = No \wedge$
$((impasse \neq 'NONE') \vee$
$(w.value \notin items) \vee$
$(number_of_winners = One \wedge$
$(\exists w_2 : WM \bullet$
$w \neq w_2 \wedge w_2.id = first(slot) \wedge$
$w_2.attribute = second(slot) \wedge$
$w_2.context_acceptable_preference = No \wedge$
$w_2.value \in items))) \bullet$
$WM' = WM \setminus \{w\}$

When no more elements need to be added to or removed from working memory, then the *WMPHaseChangeWMFinished* transition moves *WMPHase* back to the decide slot state.

<i>WMPHaseChangeWMFinished</i>
Δ <i>Decide</i>
$wmphase_state = WMPHaseChangeWMState$
$- (pre\ WMPHaseAddWME \vee pre\ WMPHaseRemoveWME)$
$wmphase_state' = WMPHasePickSlotState$

When all of the changed slots have been decided, *WMPHaseFinish* finishes *WMPHase*.

<p><i>WMPHaseFinish</i> — Δ <i>Decide</i></p> <p><i>wmphase_state</i> = <i>WMPHasePickSlotState</i></p> <p><i>changed_slots</i> \ (<i>GM</i> × <i>ContextSlotAttribute</i>) = \emptyset</p> <p><i>wmphase_state'</i> = <i>WMPHaseFinishedState</i></p>
--

The *WMPHase* is stepped by taking any one of its transitions, with the exception of the initializing transition.

$$\begin{aligned}
 \text{WMPHaseStep} \cong & \text{WMPHaseStartPickSlot} \vee \text{WMPHaseStartDecideSlot} \vee \\
 & \text{WMPHaseStepDecideSlot} \vee \text{WMPHaseStartImpasser} \vee \\
 & \text{WMPHaseStepImpasser} \vee \text{WMPHaseImpasserFinished} \vee \\
 & \text{WMPHaseAddWME} \vee \text{WMPHaseRemoveWME} \vee \\
 & \text{WMPHaseChangeWMFinished} \vee \text{WMPHaseFinish}
 \end{aligned}$$

7.10 Quiescence Phase Transitions

QPhase is very parallel in structure to *WMPHase*, but differs because it must walk the context stack instead of simply processing the changed slots. It adds in four transitions to the basic structure of working memory phase, for a total of 16 transitions (see Figure 7.4).

1. *QPhaseInitialize* — initializes *QPhase*.
2. *QPhaseReset* — resets *QPhase*'s state when its operation is abnormally terminated.
3. *QPhaseStartPickSlot* — starts *QPhase* iterating through the slots of the context stack from oldest to newest.
4. *QPhaseNextSlot* — if the slot under consideration does not need to be decided, this transition picks the next slot in the context stack to try.
5. *QPhaseNoChangedPreferencesNoChangeImpasse* — if the slot under consideration has no changed preferences and is the last slot of the context stack, this transition creates a no-change impasse.
6. *QPhaseStartDecideSlot* — this starts *QPhase* running the preference semantics.
7. *QPhaseStepDecideSlot* — this transition steps preference semantics.

8. *QPhaseStartImpasser* — when preference semantics is finished, this transition moves QPhase on to the state that handles impassing.
9. *QPhaseChangedPreferencesNoChangeImpasse* — if the slot is the last slot of the context and has changed preferences but the preference semantics does not produce an impasser or a winning candidate, then this transition creates a no-change impasse.
10. *QPhaseStepImpasser* — steps the impasser on the current slot.
11. *QPhaseImpasserFinished* — moves the QPhase from stepping the impasser into the state that changes working memory.
12. *QPhaseAddWME* — this transition adds the working memory elements for slots that preference semantics specifies.
13. *QPhaseRemoveWME* — this transition removes the working memory elements for the slot that preference semantics no longer specifies.
14. *QPhaseReAddWME* — this transition re-adds a working memory element for a slot that has been explicitly reconsidered.
15. *QPhaseSlotUnchanged* — if QPhase has not changed the slot under consideration then this loops back to the pick slot state to decide another slot.
16. *QPhaseSlotChanged* — if QPhase has changed the slot, then this ends the execution of QPhase.

QPhaseInitialize initializes the QPhase state machine by setting its state counter to the initial state.

<i>QPhaseInitialize</i> Δ Decide <hr/> <i>qphase_state'</i> = <i>QPhaseFinishedState</i> <i>qphase_state'</i> = <i>QPhaseInitialState</i>

QPhaseReset resets the QPhase's state when the machine is terminated abnormally.

<i>QPhaseReset</i> Δ Decide <hr/> <i>qphase_state'</i> = <i>QPhaseInitialState</i> <i>changed_slots</i> = \emptyset

QPhase iterates over the slots of the context from oldest slot to newest slot. *QPhaseStartPickSlot* starts the loop by setting the initial slot to the top goal and the problem-space attribute, and entering the *QPhasePickSlotState*.

QPhaseStartPickSlot

Δ *Decide*

qphase_state = *QPhaseInitialState*

slot = (*top_goal*, 'PROBLEM-SPACE')

qphase_state' = *QPhasePickSlotState*

NextContextSlot orders the slots to allow QPhase to know which slot of the goal to consider next. The relation's minimum is 'GOAL' to simplify the calculation of the slot of no-change impasses.

NextContextSlot : *ContextSlotAttribute* \rightarrow *ContextSlotAttribute*

NextContextSlot = {('GOAL', 'PROBLEM-SPACE'),
('PROBLEM-SPACE', 'STATE'),
('STATE', 'OPERATOR')}

QPhase must know when the slot it is deciding is the last one that it should consider for this impasse. If the slot is the operator, it is the last to check as it is the last in the impasse. If the context does not have a working memory element in memory for the slot, then it is the last one to check, e.g., you can't install an operator unless there is a state.

LastSlotToCheckInContext : $\mathbb{P}(\text{Slot} \times \mathbb{P} \text{WME})$

$\forall s : \text{Slot}; \text{WM} : \mathbb{P} \text{WME} \bullet$

LastSlotToCheckInContext(*s*, *WM*) \Leftrightarrow

(*second*(*s*) = 'OPERATOR' \vee

$\neg (\exists w : \text{WM} \bullet$

w.id = *first*(*s*) \wedge *w.attribute* = *second*(*s*) \wedge

w.context_acceptable_preference = *No*))

QPhase uses the *LastSlotToCheckInContextStack* predicate to determine the last one of the context stack for it to process. A slot is the last one to check in the stack if it is the last slot to check in the last impasse. It will signal a no-change impasse for the last filled slot if preference semantics can not reach a conclusion about this slot.

$$\begin{array}{l}
 \hline
 \textit{LastSlotToCheckInContextStack}_- : \mathbb{P}(\textit{Slot} \times \mathbb{P} \textit{Identifier} \times \mathbb{P} \textit{WME}) \\
 \hline
 \forall s : \textit{Slot}; GM : \mathbb{P} \textit{Identifier}; WM : \mathbb{P} \textit{WME} \bullet \\
 \textit{LastSlotToCheckInContextStack}(s, GM, WM) \Leftrightarrow \\
 \neg (\exists g : GM \bullet \textit{makeWME}(g, \textit{OBJECT}, \textit{first}(s), \textit{No}) \in WM) \wedge \\
 \textit{LastSlotToCheckInContext}(s, WM)
 \end{array}$$

QPhase must decide a context slot if there is no value for it in working memory and there are preference changes, or if there is a value in working memory and there is a reconsider.

$$\begin{array}{l}
 \hline
 \textit{DecidableSlot}_- : \mathbb{P}(\textit{Slot} \times \mathbb{P} \textit{WME} \times \mathbb{P} \textit{Preference} \times \mathbb{P} \textit{Slot}) \\
 \hline
 \forall slot : \textit{Slot}; WM : \mathbb{P} \textit{WME}; PM : \mathbb{P} \textit{Preference}; \textit{changed_slots} : \mathbb{P} \textit{Slot} \bullet \\
 \textit{DecidableSlot}(slot, WM, PM, \textit{changed_slots}) \Leftrightarrow \\
 (\neg (\exists v : \textit{Symbol} \bullet \\
 \textit{makeWME}(\textit{first}(slot), \textit{second}(slot), v, \textit{No}) \in WM) \wedge \\
 slot \in \textit{changed_slots}) \vee \\
 (\exists v : \textit{Symbol} \bullet \\
 \textit{makeWME}(\textit{first}(slot), \textit{second}(slot), v, \textit{No}) \in WM \wedge \\
 \textit{makeUnaryPreference}(\textit{first}(slot), \textit{second}(slot), v, \textit{@}) \in PM)
 \end{array}$$

While QPhase is iterating across the context stack, *QPhaseNextSlot* allows QPhase to skip slots that do not need to be decided. If the slot that QPhase is working with is not decidable and it is not the last slot to check, then it picks the next slot to decide. If the slot is not the last in this context to check, it moves to the next slot in the same context. If the slot is the last in this context to decide, it starts on the problem space of the next context.

QPhaseNextSlot

Δ Decide

qphase_state = *QPhasePickSlotState* = *qphase_state'*

\neg *DecidableSlot*(*slot*, *WM*, *PM*, *changed_slots*)

\neg *LastSlotToCheckInContextStack*(*slot*, *GM*, *WM*)

\neg *LastSlotToCheckInContext*(*slot*, *WM*) \Rightarrow
slot' = (*first*(*slot*), *NextContextSlot*(*second*(*slot*)))

LastSlotToCheckInContext(*slot*, *WM*) \Rightarrow
(\exists *sg* : *GM* : *makeWME*(*sg*, 'OBJECT', *first*(*slot*), *No*) \in *WM* •
slot' = (*sg*, 'PROBLEM-SPACE'))

If QPhase's loop down the context stack takes it to the last slot in the context, and this slot has no preference changes, this transition signals a no-change impasse. The no change impasse is for the last filled slot of the context, or if the context has no filled slot it is a goal no-change.

QPhaseNoChangedPreferencesNoChangeImpasse _____
 Δ Decide

qphase_state = *QPhasePickSlotState*
 \neg *DecidableSlot*(*slot*, *WM*, *PM*, *changed_slots*)
LastSlotToCheckInContextStack(*slot*, *GM*, *WM*)
 \neg ($\exists w : WM \bullet w.id = first(slot) \wedge w.attribute = second(slot) \wedge$
 $w.context_acceptable_preference = No) \Rightarrow$
 $slot' = (first(slot), NextContextSlot \sim (second(slot)))$
 $(\exists w : WM \bullet w.id = first(slot) \wedge w.attribute = second(slot) \wedge$
 $w.context_acceptable_preference = No) \Rightarrow$
 $slot' = (first(slot), second(slot))$
ImpasserInitialize
impasse' = `NO-CHANGE`
items' = \emptyset
qphase_state' = *QPhaseImpasserState*

If there are changed preferences for a slot, QPhase moves to the *QPhaseDecideSlotState* to step preference semantics. The slot to be decided is removed from the *changed_slots* set so that it is not decided again.

QPhaseStartDecideSlot _____
 Δ Decide

qphase_state = *QPhasePickSlotState*
DecidableSlot(*slot*, *WM*, *PM*, *changed_slots*)
 $changed_slots' = changed_slots \setminus \{slot\}$
PSInitialize
 $context_changed = Unchanged$
qphase_state' = *QPhaseDecideSlotState*

QPhaseStepDecideSlot steps preference semantics on the current slot.

<i>QPhaseStepDecideSlot</i>
Δ <i>Decide</i>
<i>qphase_state</i> = <i>QPhaseDecideSlotState</i> = <i>qphase_state'</i>
<i>PSSStep</i>

When preference semantics has finished, this transition moves *QPhase* to the *QPhaseImpasserState* to step the impasser. The impasser will add an impasse, remove an old impasse, change the items of the existing impasse, or do nothing depending upon the contents of working memory and the results of preference semantics.

<i>QPhaseStartImpasser</i>
Δ <i>Decide</i>
<i>qphase_state</i> = <i>QPhaseDecideSlotState</i>
\neg pre <i>PSSStep</i>
\neg <i>LastSlotToCheckInContextStack</i> (<i>slot</i> , <i>GM</i> , <i>WM</i>) \vee
<i>impasse</i> \neq 'NONE' \vee
<i>items</i> \neq \emptyset
<i>ImpasserInitialize</i>
<i>qphase_state'</i> = <i>QPhaseImpasserState</i>

If the slot under consideration is the last slot in the context and preference semantics has no winner, then *QPhaseChangedPreferencesNoChangeImpasse* sets the impasse to no-change before starting the impasser.

QPhaseChangedPreferencesNoChangeImpasse _____
 Δ *Decide*

qphase_state = *QPhaseDecideSlotState*
 - pre *PSStep*
items = \emptyset
impasse = `NONE`
LastSlotToCheckInContextStack(*slot*, *GM*, *WM*)
 - $(\exists w : WM \bullet w.id = first(slot) \wedge w.attribute = second(slot) \wedge w.context_acceptable_preference = No) \Rightarrow slot' = (first(slot), NextContextSlot^{\sim}(second(slot)))$
 - $(\exists w : WM \bullet w.id = first(slot) \wedge w.attribute = second(slot) \wedge w.context_acceptable_preference = No) \Rightarrow slot' = (first(slot), second(slot))$
ImpasserInitialize
qphase_state' = *QPhaseImpasserState*

QPhaseStepImpasser steps the impasser state machine.

QPhaseStepImpasser _____
 Δ *Decide*

qphase_state = *QPhaseImpasserState* = *qphase_state'*
ImpasserStep

When the impasser is finished, *QPhaseImpasserFinished* moves *QPhase* on to *ChangeWMState* to add and remove working memory elements for the slot.

QPhaseImpasserFinished _____
 Δ *Decide*

qphase_state = *QPhaseImpasserState*
 - pre *ImpasserStep*
qphase_state' = *QPhaseChangeWMState*

Whenever QPhase changes a context slot, it must clear out all of the working memory elements for the lower slots of the same context. *LaterSlotsOfContextSlot* returns the working memory elements for the later slots so that they can be removed from working memory.

$$\text{LaterSlotsOfContextSlot} : \text{Slot} \times \mathbb{P} \text{WME} \rightarrow \mathbb{P} \text{WME}$$

$$\forall s : \text{Slot}; \text{WM} : \mathbb{P} \text{WME} \bullet$$

$$\text{LaterSlotsOfContextSlot}(s, \text{WM}) =$$

$$\{w : \text{WM} \mid w.\text{id} = \text{first}(s) \wedge w.\text{context_acceptable_preference} = \text{No} \wedge \\ ((\text{second}(s) = \text{'PROBLEM-SPACE'} \wedge \\ w.\text{attribute} \in \{\text{'STATE'}, \text{'OPERATOR'}\}) \vee \\ (\text{second}(s) = \text{'STATE'} \wedge w.\text{attribute} = \text{'OPERATOR'}))\}$$

If the preference semantics produces a set of items and no impasse, and none of the items is in working memory then *QPhaseAddWME* adds a working memory element for one of the items. The item used is picked non-deterministically. It marks the context as changed so that QPhase knows that it has found a change and that it should stop iterating across the slots. As it does this, it also deletes all of the later slots for the same context.

$$\text{QPhaseAddWME}$$

$$\Delta \text{Decide}$$

$$\text{qphase_state} = \text{QPhaseChangeWMState} = \text{qphase_state}'$$

$$\text{impasse} = \text{'NONE'}$$

$$\neg (\exists w : \text{WM} \bullet w.\text{id} = \text{first}(\text{slot}) \wedge w.\text{attribute} = \text{second}(\text{slot}) \wedge \\ w.\text{context_acceptable_preference} = \text{No} \wedge w.\text{value} \in \text{items})$$

$$\text{context_changed} = \text{Unchanged} \Rightarrow \text{context_changed}' = \text{Changed}$$

$$\exists i : \text{items}; w : \text{WME} \mid w = \text{makeWME}(\text{first}(\text{slot}), \text{second}(\text{slot}), i, \text{No}) \bullet \\ (\text{GM}' = \text{GM} \setminus \text{GoalDescendants}(\text{first}(\text{slot}), \text{GM}, \text{WM}) \wedge \\ \text{WM}' = \text{WM} \setminus \text{LaterSlotsOfContextSlot}(\text{slot}, \text{WM}) \\ \cup \{w\} \wedge \\ \text{Trace WME})$$

QPhaseRemoveWME deletes any elements for the context slot that are not in the items set decided by preference semantics. It marks the context changed and deletes the elements for later slots. The one exception to this rule is that context slots are not emptied for no-change impasses. It also updates the *bottom_goal* pointer, in case it has deleted context impasses.

QPhaseRemoveWME

Δ Decide

$qphase_state = QPhaseChangeWMState = qphase_state'$

$context_changed = Unchanged \Rightarrow context_changed' = Changed$

$\exists w : WM \mid$

$w.id = first(slot) \wedge w.attribute = second(slot) \wedge$

$w.context_acceptable_preference = No \wedge$

$w.value \notin items \bullet$

$WM' = WM \setminus \{r : WM \mid r = w \wedge impasse \neq 'NO-CHANGE'\}$

$\setminus LaterSlotsOfContextSlot(slot, WM)$

$bottom_goal' = first(slot)$

Sometimes, QPhase re-decides an installed slot because it has a reconsider preference and preference semantics produces no impasse and an items set that holds the installed value. *QPhaseReAddWME* re-installs this candidate by removing the working memory element, the lower context slots, all subgoals and their architecturally created elements, and then re-adds the working memory element.

<p style="text-align: center;"><i>QPhaseReAddWME</i></p> <hr/> <p style="text-align: center;">Δ <i>Decide</i></p> <hr/> <p><i>qphase_state</i> = <i>QPhaseChangeWMState</i></p> <p><i>context_changed'</i> = <i>Changed</i></p> <p><i>impasse</i> = `NONE`</p> <p>$\exists w : WM \mid$ $w.id = first(slot) \wedge w.attribute = second(slot) \wedge$ $w.context_acceptable_preference = No \wedge$ $w.value \in items \wedge$ $makeUnaryPreference(first(slot), second(slot), w.value, '@') \in PM \bullet$ $(GM' = GM \setminus GoalDescendants(first(slot), GM, WM) \wedge$ $WM' = WM \setminus \{w\}$ $\setminus LaterSlotsOfContextSlot(slot, WM)$ $\setminus (\bigcup \{g : GoalDescendants(first(slot), GM, WM) \bullet$ $ImpassesWMEs(g, GM, WM)\})$ $\cup \{w\} \wedge$ $Trace WME)$</p> <p><i>qphase_state'</i> = <i>QPhaseFinishedState</i></p>

If working memory is already consistent with the decision for the slot then *QPhaseSlotUnchanged* will recognize that the context is unchanged and move QPhase back into the pick slot state to continue the loop over the slots.

<p style="text-align: center;"><i>QPhaseSlotUnchanged</i></p> <hr/> <p style="text-align: center;">Δ <i>Decide</i></p> <hr/> <p><i>qphase_state</i> = <i>QPhaseChangeWMState</i></p> <p>$\neg (pre\ QPhaseAddWME \vee pre\ QPhaseRemoveWME)$</p> <p><i>context_changed</i> = <i>Unchanged</i></p> <p><i>qphase_state'</i> = <i>QPhasePickSlotState</i></p>

If neither the Impasser or the *ChangeWMState* have changed memory and all adds and removes have occurred, then *QPhaseSlotChanged* finishes the execution of *QPhase*.

$$\begin{array}{l}
 \hline
 \textit{QPhaseSlotChanged} \\
 \Delta \textit{Decide} \\
 \hline
 \textit{qphase_state} = \textit{QPhaseChangeWMState} \\
 \neg (\textit{pre QPhaseAddWME} \vee \textit{pre QPhaseRemoveWME}) \\
 \textit{context_changed} = \textit{Changed} \\
 \textit{qphase_state}' = \textit{QPhaseFinishedState} \\
 \hline
 \end{array}$$

QPhaseStep drives the *QPhase* one step forward in the state machine by executing one of the *QPhase* state transitions.

$$\begin{aligned}
 \textit{QPhaseStep} \cong & \textit{QPhaseStartPickSlot} \vee \textit{QPhaseNextSlot} \vee \\
 & \textit{QPhaseNoChangedPreferencesNoChangeImpasse} \vee \\
 & \textit{QPhaseStartDecideSlot} \vee \textit{QPhaseStepDecideSlot} \vee \\
 & \textit{QPhaseStartImpasser} \vee \textit{QPhaseStepImpasser} \vee \\
 & \textit{QPhaseChangedPreferencesNoChangeImpasse} \vee \\
 & \textit{QPhaseImpasserFinished} \vee \\
 & \textit{QPhaseRemoveWME} \vee \textit{QPhaseAddWME} \vee \textit{QPhaseReAddWME} \vee \\
 & \textit{QPhaseSlotUnchanged} \vee \textit{QPhaseSlotChanged}
 \end{aligned}$$

Chapter 8

Top Level

The actions of Soar's modules are sequenced by a top level state machine that steps the various control points of each module. Recognition memory is stepped through preference phase. Decide is stepped through WMPhase and QPhase. IO is stepped through the input cycle and the output cycle.

This chapter specifies the states of the top level state machine, Soar's state schema and the transitions of the top level state machine.

8.1 The States of the Top Level

Soar's top level state machine has ten states (Figure 8.1).

1. *TLLInitialState* — the machine starts here when initialized.
2. *TLLDecisionCycleState* — the top level enters this state before starting each decision cycle.
3. *TLElaborationPhaseState* — the top level enters this state before starting each elaboration phase.
4. *TLElaborationCycleState* — the top level enters this state before starting each elaboration cycle.
5. *TLLInputCycleState* — the top level uses this state to step the input cycle state machine.
6. *TLLPreferencePhaseState* — the top level uses this state to step the preference phase state machine.
7. *TLLWorkingMemoryPhaseState* — the top level uses this state to step the working memory phase state machine.

8. *TLOutputCycleState* — the top level uses this state to step the output cycle state machine.
9. *TLDecisionPhaseState* — the top level uses this state to step the quiescence phase state machine.
10. *TLFinishedState* — the state that the top level enters when the system is halted. There is no way to reach this state in the specification because Soar is a situated agent that behaves continuously, i.e., Soar never turns itself off. However, an implementation would have commands that terminate Soar's execution by moving the top level state machine into its *TFFinishedState*.

```

TopLevelState ::= TLInitialState
                | TLDecisionCycleState
                | TLElaborationPhaseState
                | TLElaborationCycleState
                | TLInputCycleState
                | TLPreferencePhaseState
                | TLWorkingMemoryPhaseState
                | TLOutputCycleState
                | TLDecisionPhaseState
                | TLFinishedState

```

The state schema for the top level, *TL* just holds a state counter of type *TopLevelState*.

<i>TL</i>
<i>top_level_state</i> : <i>TopLevelState</i>

8.2 Soar's State Schema

The Soar state schema contains the state schemas for recognition memory, IO, and Decide. It also includes the state schema for the top level (*TL*). The *IdentifierTable* is defined to be the union of all of the component identifiers of the productions, temporary memory elements, components and traces.

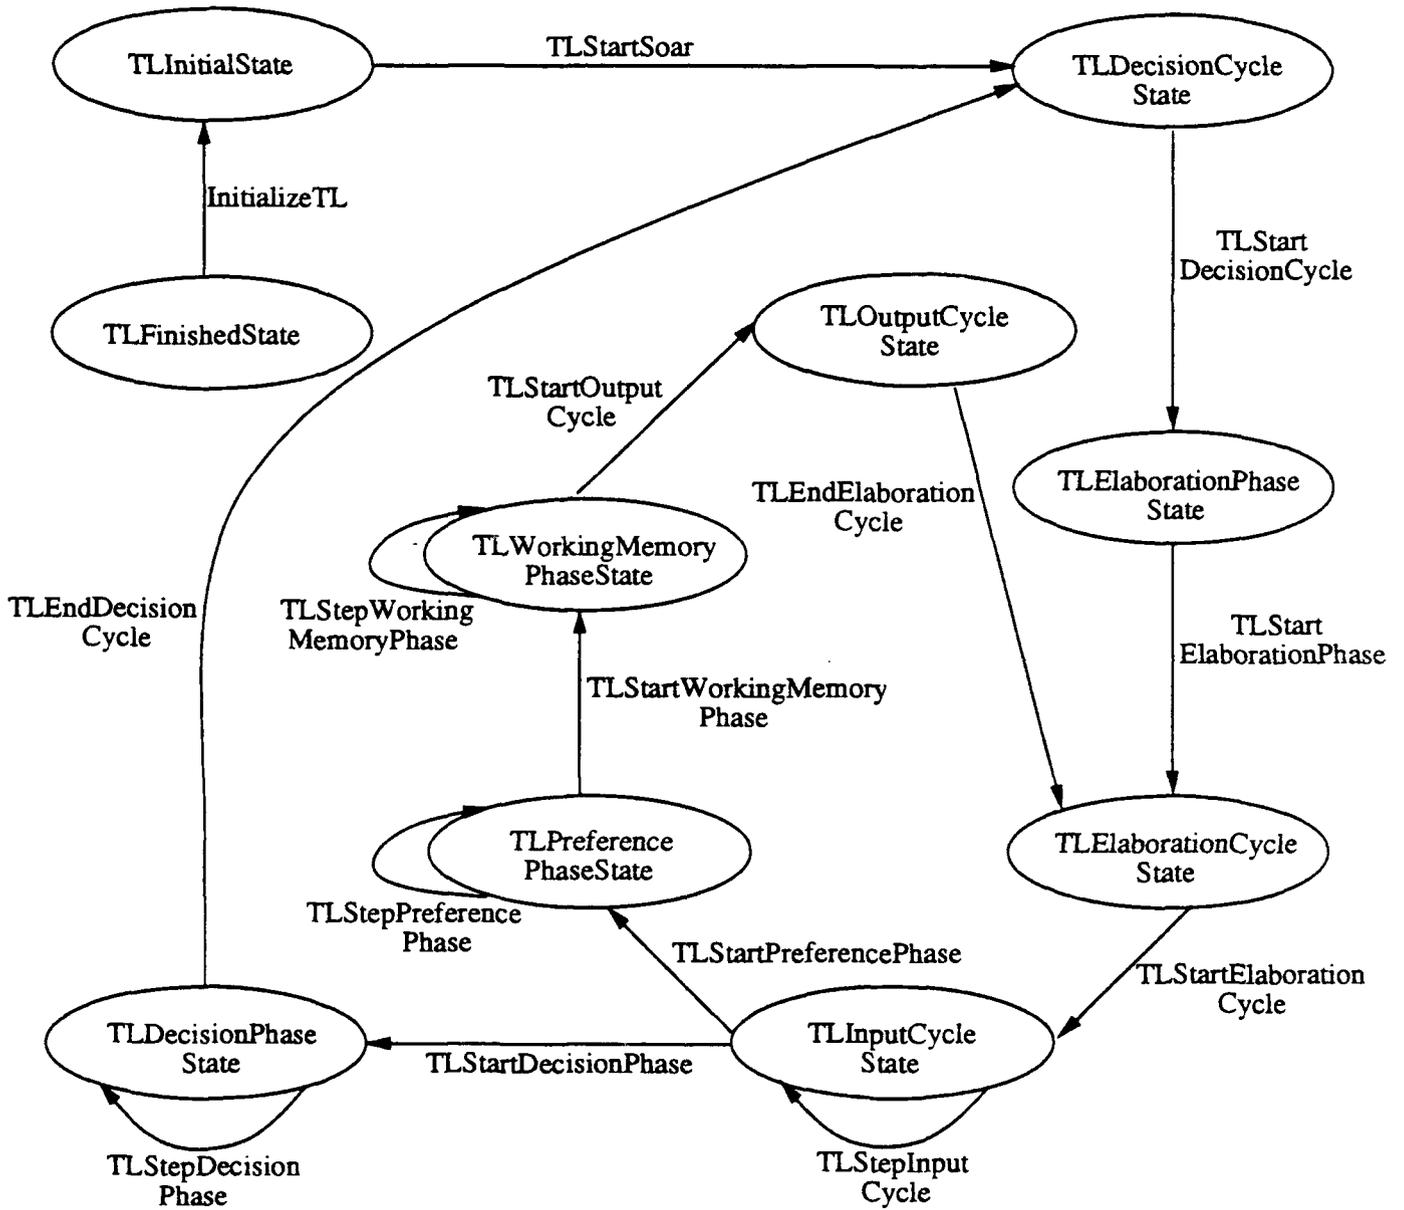


Figure 8.1: Top Level State Machine

<i>Soar</i>
<i>RM</i>
<i>IO</i>
<i>Decide</i>
<i>TL</i>
<i>IdentifierTable : P Identifier</i>
<i>IdentifierTable =</i>
$(\cup(\text{SpsComponents}(SPM) \cup$
$\text{TMEsComponents}(TM) \cup$
$\text{TracesComponents}(\text{ran TrP})) \cap$
$\text{Identifier})$

The *InitSoar* schema specifies the valid initial configurations for Soar. It constrains Soar's initial configuration to be an initial configuration for recognition memory, IO and decide, and the top level state machine must start in the *TLInitialState* state.

<i>InitSoar</i>
<i>Soar</i>
<i>InitRM</i>
<i>InitIO</i>
<i>InitDecide</i>
<i>top_level_state = TLInitialState</i>

8.3 The Top Level State Machine

TLInitialize initializes the top level state machine. Soar never initializes itself: only the user interface will ever invoke this operation.

<i>TLInitialize</i>
ΔSoar
<i>top_level_state = TLFinishedState</i>
<i>top_level_state' = TLInitialState</i>

TLReset resets the top level state machine to its initial configuration from any of its states. The top level state machine might have been stepping one or more of the hierarchy of sub-state machines, so it resets their states also.

<i>TLReset</i> Δ <i>Soar</i> <hr/> <i>top_level_state'</i> = <i>TLInitialState</i> <i>PPhaseReset</i> <i>InputCycleReset</i> <i>OutputCycleReset</i> <i>QPhaseReset</i> <i>WMPPhaseReset</i>
--

TLStartSoar moves *Soar* from its initial state into the *TLDecisionCycleState*. The *TLDecisionCycleState* is where each iteration of a decision cycle is started. When *Soar* is first started, this transition initializes decide's first goal with *DecideCreateFirstGoal*.

<i>TLStartSoar</i> Δ <i>Soar</i> <hr/> <i>top_level_state</i> = <i>TLInitialState</i> <i>DecideCreateFirstGoal</i> <hr/> <i>top_level_state'</i> = <i>TLDecisionCycleState</i>

TLStartDecisionCycle moves from the *TLDecisionCycleState* into the *TLElaborationPhaseState* starting the first elaboration phase of the decision cycle.

<i>TLStartDecisionCycle</i> Δ <i>Soar</i> <hr/> <i>top_level_state</i> = <i>TLDecisionCycleState</i> <hr/> <i>top_level_state'</i> = <i>TLElaborationPhaseState</i>

The elaboration phase consists of many elaboration cycles; so the *TLStartElaborationPhase* transition moves the top level from the *TLElaborationPhaseState* into the *TLElaborationCycleState*.

<i>TLStartElaborationPhase</i> Δ <i>Soar</i> <hr/> <i>top_level_state</i> = <i>TLElaborationPhaseState</i> <hr/> <i>top_level_state'</i> = <i>TLElaborationCycleState</i>

Each elaboration cycle is an input cycle, then a preference phase, a working memory phase and finally an output cycle. *TLStartElaborationCycle* moves from the *TLElaborationCycleState* into the *TLInputCycleState* to start the elaboration cycle by starting the input cycle. It initializes the input cycle state machine for execution, so *SoarInputCycleState* can step the input cycle state machine.

<i>TLStartElaborationCycle</i> _____ Δ <i>Soar</i> <hr/> <i>top_level_state</i> = <i>TLElaborationCycleState</i> <i>InputCycleInitialize</i> <i>top_level_state'</i> = <i>TLInputCycleState</i>
--

TLStepInputCycle moves the input cycle state machine one step.

<i>TLStepInputCycle</i> _____ Δ <i>Soar</i> <hr/> <i>top_level_state</i> = <i>TLInputCycleState</i> = <i>top_level_state'</i> <i>InputCycleStep</i>

When the input cycle state machine has finished, if *Soar* is not at *Quiescence*, *TLStartPreferencePhase* moves to *TIPreferencePhaseState* to start the preference phase.

<i>TLStartPreferencePhase</i> _____ Δ <i>Soar</i> <hr/> <i>top_level_state</i> = <i>TLInputCycleState</i> \neg pre <i>InputCycleStep</i> \neg <i>Quiescence</i> <i>top_level_state'</i> = <i>TLPreferencePhaseState</i>

If *Soar* is at *Quiescence* then *TLStartDecisionPhase* moves *Soar* to the *TLDecisionPhaseState* to start the decision, and initializes *QPhase*.

<i>TLStartDecisionPhase</i> Δ Soar <hr/> <i>top_level_state = TLInputCycleState</i> \neg pre <i>InputCycleStep</i> <i>Quiescence</i> <i>QPhaseInitialize</i> <hr/> <i>top_level_state' = TLDecisionPhaseState</i>
--

TLStepPreferencePhase steps the preference phase.

<i>TLStepPreferencePhase</i> Δ Soar <hr/> <i>top_level_state = TLPreferencePhaseState = top_level_state'</i> <hr/> <i>PPhaseStep</i>
--

When the preference phase has finished its work, *TLStartWorkingMemoryPhase* moves Soar from *TLPreferencePhaseState* into *TLWorkingMemoryPhaseState*.

<i>TLStartWorkingMemoryPhase</i> Δ Soar <hr/> <i>top_level_state = TLPreferencePhaseState</i> \neg pre <i>PPhaseStep</i> <i>WMPPhaseInitialize</i> <hr/> <i>top_level_state' = TLWorkingMemoryPhaseState</i>
--

TLStepWorkingMemoryPhase steps the working memory phase.

<i>TLStepWorkingMemoryPhase</i> Δ Soar <hr/> <i>top_level_state = TLWorkingMemoryPhaseState = top_level_state'</i> <hr/> <i>WMPPhaseStep</i>
--

When the working memory phase has finished, *TLStartOutputCycle* initializes the output cycle and moves Soar to *TLOutputCycleState*.

TLStartOutputCycle

Δ *Soar*

top_level_state = *TLWorkingMemoryPhaseState*

\neg pre *WMPhaseStep*

top_level_state' = *TLOutputCycleState*

OutputCycleInitialize

TLStepOutputCycle steps the output cycle.

TLStepOutputCycle

Δ *Soar*

top_level_state = *TLOutputCycleState* = *top_level_state'*

OutputCycleStep

When the output cycle has finished its work, *TLEndElaborationCycle* moves *Soar* back to the *ElaborationPhaseState* to start the next elaboration cycle of the elaboration phase.

TLEndElaborationCycle

Δ *Soar*

top_level_state = *TLOutputCycleState*

\neg pre *OutputCycleStep*

top_level_state' = *TLElaborationCycleState*

TLStepQPhase steps QPhase in the *DecisionPhaseState*.

TLStepQPhase

Δ *Soar*

top_level_state = *TLDecisionPhaseState* = *top_level_state'*

QPhaseStep

When the quiescence phase has finished, *TLEndDecisionCycle* moves *Soar* back to the *DecisionCycleState* to start the next decision cycle of the decision phase.

$TL_{EndDecisionCycle}$ $\Delta Soar$
$top_level_state = TL_{DecisionPhaseState}$ $\neg pre\ QPhaseStep$ $top_level_state' = TL_{DecisionCycleState}$

$TLStep$ steps the entire execution of Soar by moving the top level state machine through one of its transitions.

$$\begin{aligned}
TLStep \cong & TL_{StartSoar} \vee TL_{StartDecisionCycle} \vee \\
& TL_{StartElaborationPhase} \vee TL_{StartElaborationCycle} \vee \\
& TL_{StepInputCycle} \vee TL_{StartPreferencePhase} \vee \\
& TL_{StartDecisionPhase} \vee TL_{StepPreferencePhase} \vee \\
& TL_{StartWorkingMemoryPhase} \vee TL_{StepWorkingMemoryPhase} \vee \\
& TL_{StartOutputCycle} \vee TL_{StepOutputCycle} \vee \\
& TL_{EndElaborationCycle} \vee \\
& TL_{StepQPhase} \vee TL_{EndDecisionCycle}
\end{aligned}$$

Appendix A

Numbers

We have specified a Soar that uses only constants and identifiers to construct its temporary memory structures. Ideally, the implementation would also use exactly the same data structures. Unfortunately, Soar users find it difficult to effectively process numbers without some implementation support. The long term solution is to provide problem spaces, complete with chunk libraries, that gracefully support numeric reasoning, and some work in this direction has already been completed. In the interim, we augment our implementation's base data structures with numeric constants and extend the test structure to let productions match against them with the basic arithmetic relations.

The implementation adds in:

<i>Number</i> : \mathbb{P} <i>Constant</i>
$>, <, >=, <=$: <i>SpecialSymbol</i>
<i>NumericRelationSymbols</i> : \mathbb{P} <i>SpecialSymbol</i>
<i>NumericRelationSymbols</i> = $\{>, <, >=, <=\}$

$$Test ::= RelationalTest \langle \langle NumericRelationSymbols \times (Variable \cup Number) \rangle \rangle$$
$$\forall t : \text{ran } RelationalTest \bullet TestsComponents(t) = second(RelationalTest \sim (t))$$
$$\forall t : \text{ran } RelationalTest \bullet TestsPositiveComponents(t) = \emptyset$$

We extended $Matches_{Symbol}^{Test}$ with some extra cases to handle numeric relations and same type tests must be extended to differentiate numbers from identifiers and other constants.

$\text{Matches}_{\text{Symbol}}^{\text{Test}} - : \mathbb{P}(\text{Binding} \times \text{Test} \times \text{Symbol})$

$\forall b : \text{Binding}; t : \text{Test}; v : \text{Symbol} \bullet$

$\text{Matches}_{\text{Symbol}}^{\text{Test}}(b, t, v) \Leftrightarrow$

$((t \in \text{ran } \text{RelationalTest} \wedge$
 $(\exists \text{nrsvn} : \{\text{RelationalTest}^{\sim}(t)\} \bullet$
 $(\exists \text{vn} : \{((\text{id } \Sigma) \oplus b)(\text{second}(\text{nrsvn}))\} \bullet$
 $(\exists t : \{\text{first}(\text{nrsvn})\} \bullet$
 $(t = > \Rightarrow v > \text{vn}) \wedge$
 $(t = < \Rightarrow v < \text{vn}) \wedge$
 $(t = >= \Rightarrow v >= \text{vn}) \wedge$
 $(t = <= \Rightarrow v <= \text{vn})))))) \vee$

$(t \in \text{ran } \text{SameTypeTest} \wedge$
 $(\exists v2 : \{((\text{id } \text{Symbol}) \oplus b)(\text{SameTypeTest}^{\sim}(t))\} \bullet$
 $((v \in \text{Identifier} \wedge v2 \in \text{Identifier}) \vee$
 $(v \in \text{Constant} \setminus \text{Number} \wedge v2 \in \text{Constant} \setminus \text{Number}) \vee$
 $(v \in \text{Number} \wedge v2 \in \text{Number}))))))$

Appendix B

An Implementation Discipline

Section B.1 details a technique to allow users to observe and modify the implementation of Soar. Section B defines an implementation discipline to govern implementations of the specification. The final section, Section B.3, calculates the worst case cost of actually using the state machine representation for Soar's control.

B.1 Observing Soar's Operation

As Soar is a research vehicle for cognitive modeling and artificial intelligence work ([New90]), we would like to have an implementation that we can easily monitor and modify. We require a specification and implementation discipline that helps us to support these goals. In previous implementations, if a researcher wanted to measure the frequency of a certain type of chunk creation, she would have to read large tracts of the source code and install her statistics directly in a copy of the code. As there was no other detailed enough description of the system, only the code was an effective index into the system's operation.

Coupling this specification with an implementation discipline would allow Soar researchers to more easily solve this problem. In [PMN91] we describe a technique called *Hooking* that allows the system to execute a function, called an *Observer*, before and after the execution of any sequential operation. The technique also allows the user to, at run-time, redefine any specified operation: as long as the new implementation meets this specification of the operation in this document. Researchers could then use this specification as a guide to place observers on hooks that monitor or change the implementation's operation.

B.2 An Implementation Discipline

This section defines an implementation discipline that maps parts of this specification into executable code. This discipline was designed to allow the easy observation and modification of Soar, but the new implementation of Soar, Soar6, is not constructed using this technique.

Ideally, we would like to semi-mechanically map specifications into an optimized implementation. However, the automation of this mapping is a difficult research problem in program synthesis, so we have settled for an informal manual mapping.

The Soar specification contains four types of objects that must be mapped into the implementation:

1. state schemas — defined using schemas
2. sequential operations — defined using schemas with Δ in their signature
3. data structures — defined using base types, free types and schemas
4. state machines — built of state schemas and sequential operations

This discipline maps each state schema into a structure of the same name in the implementation. These structures contain fields for the schema's components and may contain other fields for implementation-specific information.

Mapping Sequential Operations Each sequential operation is mapped to a function of the same name in the implementation. Some of the sequential operations take inputs using the Z "?" decoration convention ([Spi89] 133) and many non-deterministically select their arguments from sets using existential quantification. The observers are passed all arguments and top level existentially quantified variables to simplify their observation of the operation.

Associated with each state changing function are two sets of observer functions, called hook sets, and an optional implementation function. The observer functions in the first, or pre-hook set, are called before the operation function executes. The observer functions in the second, or post hook set, are called after the operation function executes. The observer functions may inspect any part of Soar's state, but may not change Soar's state. The observers allow users to instrument Soar's state changes. If an implementation function has been specified it is called instead of the default sequential operation function. This would allow researchers to easily change small pieces of the implementation.

Soar's specification currently has about 150 state changing operations and most users will never care to observe or modify any of them. At the time the Soar implementation is compiled, the user would be able to turn on only the hooks she requires. Any hooks not switched on would produce no code in the resulting implementation and so have no run-time cost. If an operation has been

compiled to allow hooking, the hook set and implementation could be filled at run-time, allowing choice as to when as well as where to hook.

Mapping Data Structures The mapping to generate the data structures is much less mechanical. Implementations of data structures must be carefully selected for efficiency. Basic types would generally be represented using enumerations, structures or built-in types such as numbers. Schemas that define data types would mostly be represented using structures. Free types that define only constants would be represented as enumerated types. Free types that use injections to define types would generally be represented by structures using tags and union types. Mappings and sets would be represented in a variety of ways including pointers, lists, functions and hash tables.

Mapping State Machines The specification's state machines are constructed of state schemas and sequential operations. As above, each state schema would be represented by a structure and each transition by a function. However, many states of the machines have more than one transition leaving the state. The state machine's implementation must choose which of these transitions to apply. Every state that has more than one transition is mapped to a function of the same name that chooses which transition to apply.

Some transitions in the state machines check the pre-condition of the step operation of a sub-state machine. For example, *FinishOP* checks - pre *StepS3*. The implementation could check these pre-conditions by checking the state counter of the sub-state machine. If the machine is in any state other than the final state, its pre-condition is satisfied.

Some transitions check the pre-condition of other operations: *StartS3* checks - pre *S2*. The implementation is free to calculate this pre-condition in any way that does not allow a user of the system to observe performance different than that described by the specification. A fast way to check this is to have the state function (in this case, *S2State*'s function) record that the pre-conditions of *S2* no longer holds on *S2*'s last execution. For example, if *S2* is iterating over a set of inputs, when the set is empty it will mark that its pre-condition is no longer satisfied in a flag. When the *S2State* function is called to select the next transition, it will consult the pre-condition flag and select *StartS3*.

B.3 An Estimate of the Cost of Using State Machines

The use of state machines for control in an implementation would have a greater efficiency cost than using standard sequential code for control. This section uses an analytic and empirical analysis to demonstrate that the expense of state machine control is at worst a small percentage of Soar's total cost. First, we

describe an implementation of the state machines in detail. Second, we identify the most expensive step of state machine control. Third, we prototype just this step and benchmark it. Fourth, we estimate the number of such steps in a large Soar run and compare our measured cost with execution times from Soar5 and optimistic speed-up factors for Soar6.

The implementation of each state machines has five parts:

1. a stepping function
2. a state name counter
3. an array that the stepping function uses to index into the state functions
4. a function for each state that selects which transition to apply
5. a function for each transition

For each transition of a state machine, first the stepping function is called. It indexes using the state name into the array of state functions and calls the function for the current state. This function calculates which of the possible transitions to execute, and applies one. Thus for each step of a state machine. Soar performs three function calls and one array access. A sequential control model would represent the stepping, the array access and the state function in sequential code.

At the deepest, the specification nests state machine control only three levels. The most expensive step is likely to be the state machine to calculate preference semantics, as it must execute once for every change to any temporary memory slot. The longest execution path through this state machine has 10 transitions between 11 states plus one initialization step. This is actually an overestimate of the number of control steps required. As long as the implementation produces the same results and has equivalent observability properties, the implementation would be free to optimize the average case by exiting this sequence of steps early. For every step of preference semantics, Soar would also step the two state machines that drive it: the working memory phase and the top level.

We constructed a prototype C implementation of this control flow. It contains all of the control flow that the top level state machine and the working memory phase state machine would require to step preference semantics in a cycle of 11 transitions. On a Decstation 5000¹, the prototype ran 11 million transitions in 37 seconds. This is very likely an overestimate of an actual implementation's cost because the heavy calling of pointers to functions of this code pessimizes the use of the processor's instruction cache. When the control flow is augmented with calculation, the instruction pipeline should be able to execute some of the calculation instructions in parallel with the control flow.

The largest Soar 5.2 system that has been constructed to date runs approximately 235,000 RHS actions in 7,000 seconds on the same machine. Assuming

¹Decstation 5000 is a trade mark of the Digital Equipment Corporation.

a roughly equal rate of calls to preference semantics, Soar 5.2 requires approximately 33.6 calls to preference semantics a second, or 403.2 total control steps per second. Suppose also that we are underestimating the number of total control steps by a factor of two. This means that for every time the system takes one step inside preference semantics, it takes two steps in running the other modules of the system. Then it would require 806.4 total control steps per second of computation. Suppose very optimistically that Soar6 is 10 times faster than Soar 5.2. Soar6 would then require 8064 total state machine control steps per second. At 37 seconds per 11 million steps in Soar6, the system would spend about 2.7% of its time running the state machines.

This implementation discipline directly maps the state machines into the implementation for two reasons: they allow the system to breakpoint execution at a fine grain, and they simplify the mapping from the specification to the code. The ability to potentially breakpoint at any state in the system allows users to precisely select debugging breakpoints, stop the system, examine its current state and continue. The breakpointing will also allow the user to step the system in small increments providing a very detailed observation of the system that would facilitate learning and debugging. The small cost of at worst a few percent of total runtime seems a small cost to pay for the benefits of breakpointing.

Appendix C

Finer Grained Hooks

The granularity of the sequential operations in the specification of Soar is too coarse to allow some types of observation of the system, using the state machine implementation discipline. For example, if a researcher wanted to replace the implementation of working memory with a new data structure, she would have to replace the implementation of every operation that changed the contents of working memory. This requires re-implementing at least 13 operations.

We would like the specification to be structured to localize access to data structures to a few operations: like a program constructed using the data object module paradigm. If we were to construct sequential operations such as *AddWME*, which added an argument working memory element to working memory, we could still not easily apply them because *Z*'s sequential system model is too weak to concisely describe the common operation of iterating *AddWME* across a set of elements.

Instead of not specifying the fine grained operations, and hence not allowing their hooking, we work around this problem by requiring the implementation to use a special collection of fine grained accessors. All of the modifications to *TM*, *WM*, *PM*, *SPM*, and *IM* or *GM* will call the operations specified in this chapter. Where these sets are modified by set unions or differences, an implementation is required to repeatedly apply the individual addition and deletion operations.

C.1 TM Hooks

Temporary memory is hooked at the working memory level, the preference memory level and the temporary memory level.

C.1.1 WME Operations

AddWME adds its argument element to working memory. The element can already be in working memory, in which case it is not re-added.

<i>AddWME</i>
ΔRM
$w? : WME$
$WM' = WM \cup \{w?\}$

RemoveWME removes its argument element from working memory. The element must already exist in memory, or an implementation specific error condition is signalled.

<i>RemoveWME</i>
ΔRM
$w? : WME$
$w? \in WM$
$WM' = WM \setminus \{w?\}$

C.1.2 Preference Operations

The operation *AddPreference* adds its preference to preference memory. The preference may already exist in memory, in which case preference memory is not changed.

<i>AddPreference</i>
ΔRM
$p? : Preference$
$PM' = PM \cup \{p?\}$

The *RemovePreference* operation removes its preference from preference memory. The preference must be in memory or an implementation specific error condition is signalled.

<i>RemovePreference</i>
ΔRM
$p? : Preference$
$p? \in PM$
$PM' = PM \setminus \{p?\}$

C.1.3 TME Operations

AddTME adds its element to temporary memory. If the element is already in memory, temporary memory is not changed.

<i>AddTME</i>
ΔRM
$t? : TME$
$TM' = TM \cup \{t?\}$

RemoveTME removes its element from temporary memory. The element must be in temporary memory or an implementation specific error condition is signalled.

<i>RemoveTME</i>
ΔRM
$t? : TME$
$t? \in TM$
$TM' = TM \setminus \{t?\}$

The implementations of *AddWME*, *RemoveWME*, *AddPreference*, *RemovePreference*. *AddTME* and *RemoveTME* respect the invariant that temporary memory is constructed from working and preference memory. For example, whenever an operation calls *AddWME* the equivalent *AddTME* must also be called. Whenever an operation calls *RemoveTME* the equivalent *RemoveWME* or *RemovePreference* must also be called.

C.2 SPM Hooks

The Soar production memory, *SP*, has productions added and removed by name. The *AddSP* operation will only add in a production if there is not already one in memory of that name.

<i>AddSP</i>
ΔRM
$p? : SP$
$\neg (\exists p : SPM \bullet p.name = p?.name)$
$SPM' = SPM \cup \{p?\}$

The *RemoveSP* operation will remove a production given its name.

<i>RemoveSP</i> ΔRM <i>pname</i> : <i>Symbol</i>
$\exists p : SPM \mid p.name = pname \bullet$ $SPM' = SPM \setminus \{p\}$

C.3 GM and IM Hooks

The goal and impasse memory changes can be hooked only for context impasses, for non-context impasses or for any impasse using the six operations specified in this section.

C.3.1 Context Impasses

There is no need for an operation that hooks the addition of context impasses, because the specification contains only two operations that install them: *DecideCreateFirstGoal* and *ImpasserCreateImpasse*.

Although there is only one specified operation that removes an impasse, *ImpasserRemoveImpasse*, a user interface to Soar requires an operation to pop goals from the goal stack. *RemoveContextImpasse* is specified for this purpose. It removes a context impasse from goal memory, all of its descendent goals and all of their *ImpasseWMEs*. The implementation should behave as if a sequence of *RemoveContextImpasse* operations occur, one for each of the goals that would be removed to match this specification, from newest to oldest.

<i>RemoveContextImpasse</i> $\Delta Decide$ <i>i_or_g?</i> : <i>Identifier</i>
$i_or_g? \in GM$ $GM' = (GM \setminus \{i_or_g?\}) \setminus GoalDescendants(i_or_g?, GM, WM)$ $WM' = (WM \setminus ImpassesWMEs(i_or_g?, GM, WM))$ $\setminus (\cup \{g : GoalDescendants(i_or_g?, GM, WM) \bullet$ $ImpassesWMEs(g, GM, WM)\})$

C.3.2 Non-Context Impasses

Non-context impasses are only generated in *ImpasserCreateImpasse*, and only explicitly removed by *ImpasserRemoveImpasse*. However, some time after an impasse becomes disconnected from the context stack, it will be removed from impasse memory and it will have its working memory elements removed from

memory. For convenient hooking of this removal, *RemoveNonContextImpasse* is defined to remove the impasse from impasse memory and remove its working memory elements.

$$\begin{array}{l}
 \text{RemoveNonContextImpasse} \\
 \hline
 \Delta \text{Decide} \\
 i_or_g? : \text{Identifier} \\
 \hline
 i_or_g? \in IM \\
 IM' = IM \setminus \{i_or_g?\} \\
 WM' = WM \setminus \text{ImpassesWMEs}(i_or_g?, IM, WM) \\
 \hline
 \end{array}$$

Appendix D

The Reorderer

The reorderer is a heuristic greedy search algorithm. It searches the space of possible orderings of a production's condition elements for one that will produce a low average match cost. The reorderer does not improve the match cost of cheap chunks very much, but it is important to reduce the efficiency loss of expensive chunks.

This specification is the result of an effort to carefully understand and document the branching factor heuristic of the greedy search. The previous implementation [Sca86] lumped the conditions into five equivalence classes without carefully understanding the nature of the heuristic or the fine grained test structure of the conditions. Consequently, the calculation of the heuristic was mixed into the control structure of the search. It could not be tailored to fully take advantage of architecture and user supplied information. Worse still it threw up its hands on complicated conjunctive or simple disjunctive tests and assigned them to an expensive equivalence class.

The reorderer has three parts: a classifier, a branching factor heuristic and the greedy search itself. The classifier groups working memory element tests into classes based upon the constants and variables tested in their components and the set of bound variables and the set of variables constrained to bind to impasse identifiers. The classification is used to estimate the average match cost of the condition, called the branching factor [Tam91]. The reorderer performs a greedy search for an ordering of conditions that minimizes the branching factor of the entire ordering; with ties broken by two steps of look-ahead.

D.1 Classifying WMETests

Working memory element tests are classified bottom up: the constants of the tests are classified, then the test is classified, and finally the working memory element test is classified.

The constants and tests of the conditions are classified into seven classes:

- C— a test that tests for equality with a constant
- B— a test that tests for equality with a bound variable
- U— a test that tests for equality with an unbound variable or one that tests only for inequality with a constant or bound variable
- GIU— a test that tests for equality with an unbound goal or impasse variable
- GIB— a test that tests for equality with a bound goal or impasse variable
- YES— a test for an acceptable preference and
- NO— a test for not being an acceptable preference.

The set *TestClass* is constructed from Σ to hold the set of the possible classes.

$$\begin{array}{l} C, B, U, GIU, GIB, YES, NO : \Sigma \\ \hline TestClass : \mathcal{P} \Sigma \\ \hline TestClass = \{C, B, U, GIU, GIB, YES, NO\} \end{array}$$

ClassifySigma takes a constant or variable from a test, the set of bound variables and the set of variables bound to impasses and maps them into one of the seven classifications.

$$\begin{array}{l} \underline{ClassifySigma : \Sigma \times \mathcal{P} Variable \times \mathcal{P} Variable \rightarrow TestClass} \\ \forall vc : \Sigma; bound, impasses : \mathcal{P} Variable \bullet \\ ((vc \in Constant \Rightarrow ClassifySigma(vc, bound, impasses) = C) \wedge \\ (vc \in Variable \Rightarrow \\ (vc \in impasses \Rightarrow \\ (vc \in bound \Rightarrow \\ ClassifySigma(vc, bound, impasses) = GIB) \wedge \\ (vc \notin bound \Rightarrow \\ ClassifySigma(vc, bound, impasses) = GIU)) \wedge \\ (vc \notin impasses \Rightarrow \\ (vc \in bound \Rightarrow \\ ClassifySigma(vc, bound, impasses) = B) \wedge \\ (vc \notin bound \Rightarrow \\ ClassifySigma(vc, bound, impasses) = U)))))) \end{array}$$

The *ClassifyTest* function takes a test, the set of bound variables and the set of variables bound to impasses and classifies the test. Blank tests are classified as if they were unbound variables because, like an unbound variable, they can

match anything. Yes tests and No tests are classified into YES and NO classes. Equality tests are classified by the variable or constant that they match. Same-type tests and not tests can match against an arbitrary number of values, so they are also classified as unbound variables. Disjunctions can only match against one of a few constants, so they are classified as constants.

Conjunctions are more complicated to classify. If the conjunction contains a disjunction it can only match the constants in the disjunction, so it is classified as a constant. If it contains an equality it is assigned the classification of the equality. If it contains a not test or a same-type test, it is treated as an unbound variable. These rules are not mutually exclusive; they are intended to be applied sequentially to produce the most match restrictive classification.

ClassifyTest : *Test* × *P Variable* × *P Variable* → *TestClass*

$$\begin{aligned} &\forall t : \text{Test}; \text{bound}, \text{impasses} : \text{P Variable} \bullet \\ &((t = \text{BlankTest} \Rightarrow \text{ClassifyTest}(t, \text{bound}, \text{impasses}) = \text{U}) \wedge \\ &(t = \text{YesTest} \Rightarrow \text{ClassifyTest}(t, \text{bound}, \text{impasses}) = \text{YES}) \wedge \\ &(t = \text{NoTest} \Rightarrow \text{ClassifyTest}(t, \text{bound}, \text{impasses}) = \text{NO}) \wedge \\ &(t \in \text{ran EqualityTest} \Rightarrow \\ &\quad \text{ClassifyTest}(t, \text{bound}, \text{impasses}) = \\ &\quad \quad \text{ClassifySigma}(\text{EqualityTest}^{\sim}(t), \text{bound}, \text{impasses})) \wedge \\ &(t \in \text{ran NotTest} \Rightarrow \text{ClassifyTest}(t, \text{bound}, \text{impasses}) = \text{U}) \wedge \\ &(t \in \text{ran SameTypeTest} \Rightarrow \text{ClassifyTest}(t, \text{bound}, \text{impasses}) = \text{U}) \wedge \\ &(t \in \text{ran DisjunctiveTest} \Rightarrow \text{ClassifyTest}(t, \text{bound}, \text{impasses}) = \text{C}) \wedge \\ &(t \in \text{ran ConjunctiveTest} \Rightarrow \\ &((\exists c : \text{ConjunctiveTest}^{\sim}(t) \bullet c \in \text{ran DisjunctiveTest}) \Rightarrow \\ &\quad \text{ClassifyTest}(t, \text{bound}, \text{impasses}) = \text{C}) \wedge \\ &((\exists c : \text{ConjunctiveTest}^{\sim}(t) \bullet c \in \text{ran EqualityTest}) \Rightarrow \\ &\quad (\exists c : \text{ConjunctiveTest}^{\sim}(t) \cap \text{ran EqualityTest} \bullet \\ &\quad \quad \text{ClassifyTest}(t, \text{bound}, \text{impasses}) = \\ &\quad \quad \quad \text{ClassifySigma}(\text{EqualityTest}^{\sim}(c), \text{bound}, \text{impasses}))) \wedge \\ &((\exists c : \text{ConjunctiveTest}^{\sim}(t) \bullet c \in \text{ran NotTest}) \\ &\quad \Rightarrow \text{ClassifyTest}(t, \text{bound}, \text{impasses}) = \text{U}) \wedge \\ &((\exists c : \text{ConjunctiveTest}^{\sim}(t) \bullet c \in \text{ran SameTypeTest}) \Rightarrow \\ &\quad \text{ClassifyTest}(t, \text{bound}, \text{impasses}) = \text{U}))) \end{aligned}$$

The *WMETestClass* schema holds classifications for an entire *WMEtest*. The *WME* test class uses Σ as the value of the attribute and value components so that the classes can be used to map attribute and value specific classifications to branching factors. The system and the user can then enter default information about the branching factors of conditions in a uniform manner.

WMETestClass

id : *TestClass*
attribute : Σ
value : Σ
context_acceptable_preference : {YES, NO}

The description of the derivation of the branching factor table requires the instantiation of 90 *WME* test classes. For brevity, the *WC* constructor is defined to instantiate the *WME* test classes.

WC : *TestClass* \times *TestClass* \times *TestClass* \times *TestClass* \rightarrow *WMETestClass*

\forall *id* : *TestClass*; *attribute* : *TestClass*; *value* : *TestClass*;
context_acceptable_preference : {YES, NO} •
WC(*id*, *attribute*, *value*, *context_acceptable_preference*) =
(μ *id* : {*id*}; *attribute* : {*attribute*}; *value* : {*value*};
context_acceptable_preference : {*context_acceptable_preference*} •
 θ *WMETestClass*)

ClassifyWMEtest classifies a *WME* test given the set of bound variables and the set of variables bound to impasses. For the attribute and values the goal and impasse variables are not passed to *ClassifyTest* so that classify test will classify goal variables in these slots as normal variables.

ClassifyWMEtest : *WMEtest* \times *P Variable* \times *P Variable* \rightarrow *WMETestClass*

\forall *wt* : *WMEtest*; *bound*, *impasses* : *P Variable* •
ClassifyWMEtest(*wt*, *bound*, *impasses*) =
WC(*ClassifyTest*(*wt.id*, *bound*, *impasses*),
ClassifyTest(*wt.attribute*, *bound*, \emptyset),
ClassifyTest(*wt.value*, *bound*, \emptyset),
ClassifyTest(*wt.context_acceptable_preference*, \emptyset , \emptyset))

D.2 Deriving the Branching Factor Heuristic

The branching factor is a measure of the average cost to match a single condition, given the set of bound variables and the set of variables bound to impasses. The principled derivation of this heuristic is complicated. There are 90 ($5 * 3 * 3 *$

2) equivalence classes of WME test; as the identifier slot of a WME test can be classified five ways, the attribute and value each three ways and the preference two ways.

The analysis of these ninety cases into branching factors is derived from five measures of the average branching of the graph of working memory.

1. G — the average number of goal and impasses in the system.
2. O — the average number of objects in the system.
3. A — the average number of attributes per object in the system.
4. V — the average number of values for any attribute in the system.
5. AP — the average number of context acceptable preferences per values.

$$\frac{G, O, A, V, AP : Z}{G > 0 \wedge O > 0 \wedge A > 0 \wedge V > 0 \wedge AP > 0}$$

The branching factor for each WME test is calculated by multiplying together the average measure for each slot that the condition leaves unconstrained.

The following table derives the average case match estimate for all 90 cases. This table is provided only to allow discussion of the correctness of this measure. The implementation will use a closed form of this function that is specified next.

1. $WC(C,C,C,YES) = 1$
2. $WC(C,C,C,NO) = 1$
3. $WC(C,C,B,YES) = 1$
4. $WC(C,C,B,NO) = 1$
5. $WC(C,C,U,YES) = V * AP$
6. $WC(C,C,U,NO) = V$
7. $WC(C,B,C,YES) = 1$
8. $WC(C,B,C,NO) = 1$
9. $WC(C,B,B,YES) = 1$
10. $WC(C,B,B,NO) = 1$
11. $WC(C,B,U,YES) = V * AP$
12. $WC(C,B,U,NO) = V$

13. $WC(C,U,C,YES) = A$
14. $WC(C,U,C,NO) = A$
15. $WC(C,U,B,YES) = A$
16. $WC(C,U,B,NO) = A$
17. $WC(C,U,U,YES) = A * V * AP$
18. $WC(C,U,U,NO) = A * V$
19. $WC(B,C,C,YES) = 1$
20. $WC(B,C,C,NO) = 1$
21. $WC(B,C,B,YES) = 1$
22. $WC(B,C,B,NO) = 1$
23. $WC(B,C,U,YES) = V * AP$
24. $WC(B,C,U,NO) = V$
25. $WC(B,B,C,YES) = 1$
26. $WC(B,B,C,NO) = 1$
27. $WC(B,B,B,YES) = 1$
28. $WC(B,B,B,NO) = 1$
29. $WC(B,B,U,YES) = V * AP$
30. $WC(B,B,U,NO) = V$
31. $WC(B,U,C,YES) = A$
32. $WC(B,U,C,NO) = A$
33. $WC(B,U,B,YES) = A$
34. $WC(B,U,B,NO) = A$
35. $WC(B,U,U,YES) = A * V * AP$
36. $WC(B,U,U,NO) = A * V$
37. $WC(U,C,C,YES) = O$
38. $WC(U,C,C,NO) = O$
39. $WC(U,C,B,YES) = O$

40. $WC(U,C,B,NO) = O$
41. $WC(U,C,U,YES) = O * V * AP$
42. $WC(U,C,U,NO) = O * V$
43. $WC(U,B,C,YES) = O$
44. $WC(U,B,C,NO) = O$
45. $WC(U,B,B,YES) = O$
46. $WC(U,B,B,NO) = O$
47. $WC(U,B,U,YES) = O * V * AP$
48. $WC(U,B,U,NO) = O * V$
49. $WC(U,U,C,YES) = O * A$
50. $WC(U,U,C,NO) = O * A$
51. $WC(U,U,B,YES) = O * A$
52. $WC(U,U,B,NO) = O * A$
53. $WC(U,U,U,YES) = O * A * V * AP$
54. $WC(U,U,U,NO) = O * A * V$
55. $WC(GIU,C,C,YES) = G$
56. $WC(GIU,C,C,NO) = G$
57. $WC(GIU,C,B,YES) = G$
58. $WC(GIU,C,B,NO) = G$
59. $WC(GIU,C,U,YES) = G * V * AP$
60. $WC(GIU,C,U,NO) = G * V$
61. $WC(GIU,B,C,YES) = G$
62. $WC(GIU,B,C,NO) = G$
63. $WC(GIU,B,B,YES) = G$
64. $WC(GIU,B,B,NO) = G$
65. $WC(GIU,B,U,YES) = G * V * AP$
66. $WC(GIU,B,U,NO) = G$

- 67. $WC(GIU, U, C, YES) = G * A$
- 68. $WC(GIU, U, C, NO) = G * A$
- 69. $WC(GIU, U, B, YES) = G * A$
- 70. $WC(GIU, U, B, NO) = G * A$
- 71. $WC(GIU, U, U, YES) = G * A * V * AP$
- 72. $WC(GIU, U, U, NO) = G * A * V$
- 73. $WC(GIB, C, C, YES) = 1$
- 74. $WC(GIB, C, C, NO) = 1$
- 75. $WC(GIB, C, B, YES) = 1$
- 76. $WC(GIB, C, B, NO) = 1$
- 77. $WC(GIB, C, U, YES) = V * AP$
- 78. $WC(GIB, C, U, NO) = V$
- 79. $WC(GIB, B, C, YES) = 1$
- 80. $WC(GIB, B, C, NO) = 1$
- 81. $WC(GIB, B, B, YES) = 1$
- 82. $WC(GIB, B, B, NO) = 1$
- 83. $WC(GIB, B, U, YES) = V * AP$
- 84. $WC(GIB, B, U, NO) = V$
- 85. $WC(GIB, U, C, YES) = A$
- 86. $WC(GIB, U, C, NO) = A$
- 87. $WC(GIB, U, B, YES) = A$
- 88. $WC(GIB, U, B, NO) = A$
- 89. $WC(GIB, U, U, YES) = A * V * AP$
- 90. $WC(GIB, U, U, NO) = A * V$

This formulation is overly-general because each production's conditions must all be connected to an impasse. So all of the cases of the form $WC(U, \dots)$ will never be used. The current system assumes that any conditions of the form: $WC(B, C, U, NO)$ are uni-attributes, e.g., have a branching factor of 1, unless told otherwise with a multi-attributes declaration.

The branching factor table of 90 cases is too large for easy use. A closed form for this table, named ABF , can be derived from the observation that the average branching factor of a WME test class is the average branching factor of its identifier ($ABFI$) multiplied by the average branching factor of its attribute ($ABFA$) and the average branching factor of its value acceptable preference combination ($ABFVAP$).

$$ABF, ABFI, ABFA, ABFVAP : WMETestClass \rightarrow Z$$

$$\forall wtc : WMETestClass \bullet$$

$$ABF(wtc) = ABFI(wtc) * ABFA(wtc) * ABFVAP(wtc)$$

$$\forall wtc : WMETestClass \bullet$$

$$((wtc.id \in \{C, B, GIB\} \Rightarrow ABFI(wtc) = 1) \wedge \\ (wtc.id = GIU \Rightarrow ABFI(wtc) = G) \wedge \\ (wtc.id = U \Rightarrow ABFI(wtc) = O))$$

$$\forall wtc : WMETestClass \bullet$$

$$((wtc.attribute = U \Rightarrow ABFA(wtc) = A) \wedge \\ (wtc.attribute \neq U \Rightarrow ABFA(wtc) = 1))$$

$$\forall wtc : WMETestClass \bullet$$

$$((wtc.value \neq U \Rightarrow ABFVAP(wtc) = 1) \wedge \\ (wtc.value = U \wedge wtc.context_acceptable_preference = YES \\ \Rightarrow ABFVAP(wtc) = V * AP) \wedge \\ (wtc.value = U \wedge wtc.context_acceptable_preference = NO \\ \Rightarrow ABFVAP(wtc) = V))$$

Soar users may care to provide information about branching factors of specific condition classes, and the system should have default knowledge about conditions that match context stack attributes. The user only cares to provide an average branching factor for the multi-attribute cases:

$WC(B, C, U, NO)$ and $WC(B, C, U, YES)$.

Using this classification scheme, the system can take advantage of some knowledge of the average case structure of the context stack:

- $WC(GIU, 'OBJECT', 'NIL', NO) = 1$ and $WC(GIB, 'OBJECT', 'NIL', NO) = 1$ — there is only one goal with no supergoal.
- $WC(GIU, 'OBJECT', U, NO) = 1$ and $WC(GIB, 'OBJECT', U, NO) = 1$ — there is only one supergoal for each

goal.

- $WC(GIU, \text{'PROBLEM-SPACE'}, U, NO) = G$,
 $WC(GIU, \text{'STATE'}, U, NO) = G$ and
 $WC(GIU, \text{'OPERATOR'}, U, NO) = G$ — on average every goal has a problem space, state and operator.
- $WC(GIB, \text{'PROBLEM-SPACE'}, U, NO) = 1$,
 $WC(GIB, \text{'STATE'}, U, NO) = 1$ and
 $WC(GIB, \text{'OPERATOR'}, U, NO) = G$ — each specific goal has on average one problem-space, state or operator.
- $WC(GIB, \text{'ITEM'}, U, NO) > 1$ — on average the items augmentation of a goal holds more than one thing.

The user branching factor function, UBF , is provided to capture the user's knowledge of condition branching factors. A user interface, like Soar5's multi-attributes ([LCAS90] 178), should be provided to allow the user to extend UBF .

$$UBF : WMETestClass \rightarrow Z$$

The system stores its knowledge of branching factors in the default branching factor function, DBF . DBF assumes that the $WC(B, C, U, NO)$ case is a uni-attribute.

$$DBF : WMETestClass \rightarrow Z$$

$$DBF = \{ (WC(GIU, \text{'OBJECT'}, \text{'NIL'}, NO), 1), \\ (WC(GIB, \text{'OBJECT'}, \text{'NIL'}, NO), 1), \\ (WC(GIU, \text{'OBJECT'}, U, NO), 1), \\ (WC(GIB, \text{'OBJECT'}, U, NO), 1), \\ (WC(GIU, \text{'PROBLEM-SPACE'}, U, NO), G), \\ (WC(GIU, \text{'STATE'}, U, NO), G), \\ (WC(GIU, \text{'OPERATOR'}, U, NO), G), \\ (WC(GIB, \text{'PROBLEM-SPACE'}, U, NO), 1), \\ (WC(GIB, \text{'STATE'}, U, NO), 1), \\ (WC(GIB, \text{'OPERATOR'}, U, NO), 1), \\ (WC(GIB, \text{'ITEM'}, U, NO), 2), \\ (WC(B, C, U, NO), 1) \}$$

BF takes two WME tests for arguments. The first is the WME test generated by the classifier. The second is the classifier with one of the possible constants of the attribute and/or value tests installed instead. *BF* checks if the system or the user has any constant specific branching information. If so it is used, otherwise the system uses the average branching factor.

$$\underline{BF : WMETestClass \times WMETestClass \rightarrow Z}$$

$$\begin{aligned} \forall \text{prototype, instantiated} : WMETestClass \bullet \\ (\text{instantiated} \in \text{dom}(DBF \oplus UBF) \Rightarrow \\ BF(\text{prototype, instantiated}) = (DBF \oplus UBF)(\text{instantiated})) \wedge \\ (\text{instantiated} \notin \text{dom}(DBF \oplus UBF) \Rightarrow \\ BF(\text{prototype, instantiated}) = ABF(\text{prototype})) \end{aligned}$$

The reorderer must check the constant specific branching factor information for any positive constants that appear in attribute or value tests.

$$\underline{ContainsPositiveConstants_ : \mathcal{P} \text{ Test}}$$

$$\begin{aligned} \forall t : \text{Test} \bullet \\ ContainsPositiveConstants(t) \Leftrightarrow \\ (t \in \text{ran DisjunctiveTest} \vee \\ t \in \text{ran EqualityTest} \wedge EqualityTest^{\sim}(t) \in \text{Constant} \vee \\ t \in \text{ran ConjunctiveTest} \wedge \\ (\exists c : ConjunctiveTest^{\sim}(t) \bullet c \in \text{ran DisjunctiveTest} \vee \\ (c \in \text{ran EqualityTest} \wedge EqualityTest^{\sim}(c) \in \text{Constant}))) \end{aligned}$$

The Soar5 reorderer would not correctly classify many conditions that involved disjunctions or conjunctions in the attributes. *SBF* carefully picks the maximum branching factor for the set of constants of the attribute or value tests of the *WME* test.

$$\begin{aligned}
 & \text{SBF} : \text{WME}TestClass \times \text{WME}Test \rightarrow \mathbb{Z} \\
 & \forall wtc : \text{WME}TestClass; wt : \text{WME}Test \bullet \\
 & \quad ((\text{ContainsPositiveConstants}(wt.attribute) \wedge \\
 & \quad \quad \text{ContainsPositiveConstants}(wt.value) \Rightarrow \\
 & \quad \quad \text{SBF}(wtc, wt) = \\
 & \quad \quad \quad \max\{a : \text{TestsPositiveComponents}(wt.attribute) \cap \text{Constant}; \\
 & \quad \quad \quad v : \text{TestsPositiveComponents}(wt.value) \cap \text{Constant} \bullet \\
 & \quad \quad \quad \text{BF}(wtc, WC(wtc.id, a, v, wtc.context_acceptable_preference))\}) \wedge \\
 & \quad (\text{ContainsPositiveConstants}(wt.attribute) \wedge \\
 & \quad \quad \neg \text{ContainsPositiveConstants}(wt.value) \Rightarrow \\
 & \quad \quad \text{SBF}(wtc, wt) = \\
 & \quad \quad \quad \max\{a : \text{TestsPositiveComponents}(wt.attribute) \cap \text{Constant} \bullet \\
 & \quad \quad \quad \text{BF}(wtc, WC(wtc.id, a, wtc.value, wtc.context_acceptable_preference))\}) \wedge \\
 & \quad (\neg \text{ContainsPositiveConstants}(wt.attribute) \wedge \\
 & \quad \quad \neg \text{ContainsPositiveConstants}(wt.value) \Rightarrow \\
 & \quad \quad \text{SBF}(wtc, wt) = \text{ABF}(wtc) = \text{BF}(wtc, wtc)))
 \end{aligned}$$

SBF will allow the reorderer to correctly order conditions, like this one from the default productions, containing the common idiom of binding to a disjunction.

```

(<s> ^{ << required-success success partial-success
... >> <svalue> } <eb> )

```

In this case, the reorderer can tell that this disjunction is a uni-attribute, not a multi-attribute.

The concept of a WME test's branching factor can be elevated to conditions. The branching factor of a positive condition is the branching factor of its WME test. Negative conditions can not be added until all of the variables that they share with the positive conditions have been bound. When a negative condition is added, its branching factor is one because it only filters the set of matches. For convenience, if the negative condition does not have all of its variables bound it is assigned an arbitrarily huge branching factor.

$$\begin{array}{l}
\text{ConditionsBranchingFactor :} \\
\text{Condition} \times \mathbb{P} \text{ Variable} \times \mathbb{P} \text{ Variable} \times \mathbb{P} \text{ Variable} \rightarrow \mathbb{Z} \\
\hline
\forall c : \text{Condition}; \text{bound, positive, impasses} : \mathbb{P} \text{ Variable} \bullet \\
((c \in \text{ran PositiveCondition} \Rightarrow \\
(\exists wt : \text{WMETest} \mid wt = \text{PositiveCondition} \sim (c) \bullet \\
\text{ConditionsBranchingFactor}(c, \text{bound}, \text{positive}, \text{impasses}) = \\
\text{SBF}(\text{ClassifyWMETest}(wt, \text{bound}, \text{impasses}), wt))) \wedge \\
(c \in \text{ran NegativeCondition} \cup \text{ran NegativeConjunctiveCondition} \Rightarrow \\
(((\text{ConditionsPositiveComponents}(c) \cap \text{Variable}) \cap \text{positive} \subseteq \text{bound}) \Rightarrow \\
\text{ConditionsBranchingFactor}(c, \text{bound}, \text{positive}, \text{impasses}) = 1) \wedge \\
(\neg ((\text{ConditionsPositiveComponents}(c) \cap \text{Variable}) \cap \text{positive} \subseteq \text{bound}) \Rightarrow \\
\text{ConditionsBranchingFactor}(c, \text{bound}, \text{positive}, \text{impasses}) = \text{Infinity})))
\end{array}$$

D.3 The Reorderer's State

The reorderer's state machine uses three states:

1. *ReordererInitialState* — the machine starts in this state.
2. *ReordererAddConditionState* — the machine adds conditions in this state.
3. *ReordererFinishedState* — the machine finishes in this state.

$$\begin{array}{l}
\text{ReordererState} ::= \text{ReordererInitialState} \\
\quad \mid \text{ReordererAddConditionState} \\
\quad \mid \text{ReordererFinishedState}
\end{array}$$

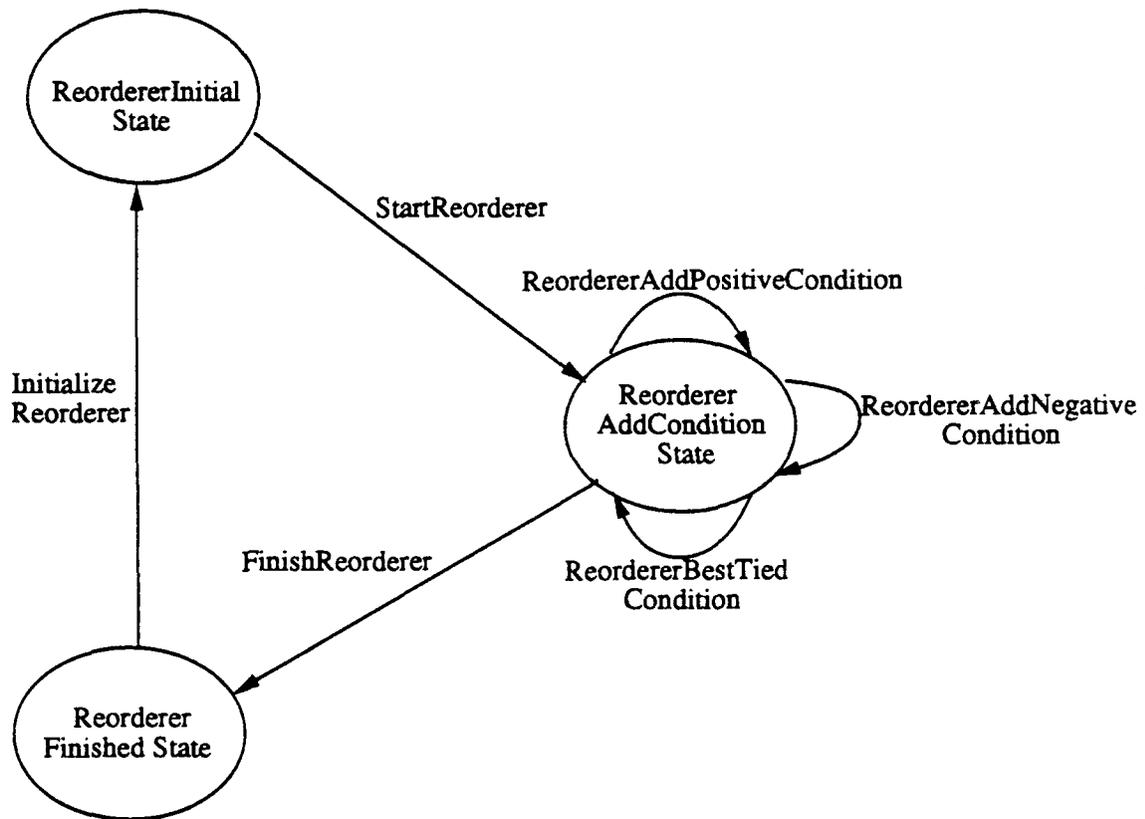


Figure D.1: Reorderer State Machine

The *Reorderer* schema holds a state counter and seven pieces of state:

1. *ordering* — the sequence of conditions selected so far
2. *conditions* — the set of conditions of the production's LHS that have not been added to the ordering
3. *impasses* — the set of variables that the condition binds to goal and impasse identifiers
4. *bound* — the set of variables bound in a condition already in the ordering
5. *positive* — the set of variables bound in all of the positive conditions and
6. *branching_factor* — the estimated average cost of matching the entire ordering.

Reorderer

reorderer_state : *ReordererState*
ordering : seq *Condition*
conditions : P *Condition*
impasses, bound, positive : P *Variable*
branching_factor : \mathbb{N}_1

When the reorderer is initialized it must start with a set of conditions to order and an empty ordering. The branching factor of an empty ordering is defined to be one. The set of bound variables is empty because no conditions have been added to the ordering. Only goal or impasse conditions can be selected to start, so there must be some goal or impasse variables. The positive variables are the variables bound positively in the positive conditions.

InitReorderer

Reorderer

conditions $\neq \emptyset$
ordering = $\langle \rangle$
bound = \emptyset
branching_factor = 1
impasses $\neq \emptyset$
positive =
 $(\bigcup (\text{ConditionsPositiveComponents} (\text{conditions} \cap \text{ran PositiveCondition})))$
 $\cap \text{Variable}$

When the reorderer's state is changed the root conditions and positive variables remain unchanged. With each step, the set of bound conditions grows and the branching factor is non-decreasing. The last predicate ensures that the reorderer either adds a condition to the end of ordering or leaves it unchanged.

Δ Reorderer
Reorderer
Reorderer'
$impasses' = impasses$
$positive' = positive$
$bound \subseteq bound'$
$branching_factor' \geq branching_factor$
$ordering' = ordering \vee tail(ordering) = ordering'$

D.4 The Reorderer's Operations

InitializeReorderer moves the state machine from the finished state into the initial state and ensures that the reorderer's configuration is a valid initial configuration.

<i>InitializeReorderer</i>
Δ Reorderer
$reorderer_state = ReordererFinishedState$
<i>InitReorderer</i>
$reorderer_state' = ReordererInitialState$

The machine's first action is to move to the add condition state.

<i>StartReorderer</i>
Δ Reorderer
$reorderer_state = ReordererInitialState$
$reorderer_state' = ReordererAddConditionState$

The *BestConditions* function selects the subset of conditions that have a minimum branching factor.

$$\begin{array}{l}
 \text{BestConditions :} \\
 \text{P Condition} \times \text{P Variable} \times \text{P Variable} \times \text{P Variable} \rightarrow \text{P Condition} \\
 \hline
 \forall C : \text{P Condition}; \text{bound, positive, impasses} : \text{P Variable} \bullet \\
 \text{BestConditions}(C, \text{bound}, \text{positive}, \text{impasses}) = \\
 \{c : C \mid \\
 \quad - (\exists c_2 : C \bullet \\
 \quad \quad \text{ConditionsBranchingFactor}(c_2, \text{bound}, \text{positive}, \text{impasses}) < \\
 \quad \quad \text{ConditionsBranchingFactor}(c, \text{bound}, \text{positive}, \text{impasses}))\}
 \end{array}$$

When there is a unique best condition or the branching factor of all of the conditions is one, the reorderer will add a condition from the set of best conditions. The reorderer picks any one of the best conditions and adds it to the end of the ordering. It updates the estimate of the cost of the ordering by multiplying it by the condition's branching factor. The positive variables of the added condition are added into the set of bound variables.

$$\begin{array}{l}
 \text{ReordererAddBestCondition} \\
 \hline
 \Delta \text{Reorderer} \\
 \hline
 \text{reorderer_state} = \text{ReordererAddConditionState} = \text{reorderer_state}' \\
 \#(\text{BestConditions}(\text{conditions}, \text{bound}, \text{positive}, \text{impasses})) = 1 \vee \\
 (\forall c : \text{BestConditions}(\text{conditions}, \text{bound}, \text{positive}, \text{impasses}) \bullet \\
 \quad \text{ConditionsBranchingFactor}(c, \text{bound}, \text{positive}, \text{impasses}) = 1) \\
 \exists c : \text{BestConditions}(\text{conditions}, \text{bound}, \text{positive}, \text{impasses}) \bullet \\
 (\text{conditions}' = \text{conditions} \setminus \{c\} \wedge \\
 \text{ordering}' = \text{ordering} \hat{\ } \langle c \rangle \wedge \\
 \text{branching_factor}' = \text{branching_factor} * \\
 \quad \text{ConditionsBranchingFactor}(c, \text{bound}, \text{positive}, \text{impasses}) \wedge \\
 \text{bound}' = \text{bound} \cup (\text{ConditionsPositiveComponents}(c) \cap \text{Variable}))
 \end{array}$$

LookAhead is given: a condition, the set of all conditions yet to be chosen, the bound variables, the positive variables and the goal or impasse variables. It calculates the branching factor for first adding the condition and then adding the best of the remaining conditions.

LookAhead :

$$\text{Condition} \times \mathbb{P} \text{Condition} \times \mathbb{P} \text{Variable} \times \mathbb{P} \text{Variable} \times \mathbb{P} \text{Variable} \rightarrow \mathbb{Z}$$

$\forall c_1 : \text{Condition}; C : \mathbb{P} \text{Condition}; \text{bound}, \text{positive}, \text{impasses} : \mathbb{P} \text{Variable} \quad c_1 \in C \bullet$
 $\exists \text{newbounds} : \mathbb{P} \text{Variable} :$
 $\text{newbounds} = \text{bound} \cup (\text{ConditionsPositiveComponents}(c_1) \cap \text{Variable}) \bullet$
 $\exists c_2 : \text{BestConditions}(C \setminus \{c_1\}, \text{bound}, \text{positive}, \text{impasses}) \bullet$
 $\text{LookAhead}(c_1, C, \text{bound}, \text{positive}, \text{impasses}) =$
 $\text{ConditionsBranchingFactor}(c_1, \text{bound}, \text{positive}, \text{impasses}) *$
 $\text{ConditionsBranchingFactor}(c_2, \text{newbounds}, \text{positive}, \text{impasses})$

BestTiedConditions is used to break ties between conditions. It returns the set of conditions that have the minimum branching factor after one step of lookahead.

BestTiedConditions :

$$\mathbb{P} \text{Condition} \times \mathbb{P} \text{Variable} \times \mathbb{P} \text{Variable} \times \mathbb{P} \text{Variable} \rightarrow \mathbb{P} \text{Condition}$$

$\forall C : \mathbb{P} \text{Condition}; \text{bound}, \text{positive}, \text{impasses} : \mathbb{P} \text{Variable} \bullet$
 $\text{BestTiedConditions}(C, \text{bound}, \text{positive}, \text{impasses}) =$
 $\{c_1 : \text{BestConditions}(C, \text{bound}, \text{positive}, \text{impasses})$
 $\quad - (\exists c_2 : \text{BestConditions}(C, \text{bound}, \text{positive}, \text{impasses}) \bullet$
 $\quad \quad \text{LookAhead}(c_2, C \setminus \{c_1\}, \text{bound}, \text{positive}, \text{impasses}) <$
 $\quad \quad \text{LookAhead}(c_1, C \setminus \{c_1\}, \text{bound}, \text{positive}, \text{impasses}))\}$

Whenever there is a tie for the best condition and all of the tied conditions have a branching factor of greater than one, *ReordererAddBestTiedCondition* selects one of the best tied conditions and adds it to the ordering.

$$\begin{array}{l}
 \text{--- } \underline{\text{ReordererAddBestTiedCondition}} \text{ ---} \\
 \Delta \text{Reorderer} \\
 \text{---} \\
 \text{reorderer_state} = \text{ReordererAddConditionState} = \text{reorderer_state}' \\
 \neg (\#(\text{BestConditions}(\text{conditions}, \text{bound}, \text{positive}, \text{impasses})) = 1) \vee \\
 (\exists c : \text{BestConditions}(\text{conditions}, \text{bound}, \text{positive}, \text{impasses}) \bullet \\
 \quad \text{ConditionsBranchingFactor}(c, \text{bound}, \text{positive}, \text{impasses}) = 1) \\
 \exists c : \text{BestTiedConditions}(\text{conditions}, \text{bound}, \text{positive}, \text{impasses}) \bullet \\
 \quad (\text{conditions}' = \text{conditions} \setminus \{c\} \wedge \\
 \quad \text{ordering}' = \text{ordering} \hat{\ } (c) \wedge \\
 \quad \text{branching_factor}' = \\
 \quad \quad \text{branching_factor} * \\
 \quad \quad \text{ConditionsBranchingFactor}(c, \text{bound}, \text{positive}, \text{impasses}) \wedge \\
 \quad \text{bound}' = \text{bound} \cup (\text{ConditionsPositiveComponents}(c) \cap \text{Variable})) \\
 \text{---}
 \end{array}$$

When all of the conditions have been added to the ordering, *FinishReorderer* moves the reorderer into the finished state.

$$\begin{array}{l}
 \text{--- } \underline{\text{FinishReorderer}} \text{ ---} \\
 \Delta \text{Reorderer} \\
 \text{---} \\
 \text{reorderer_state} = \text{ReordererAddConditionState} \\
 \text{conditions} = \emptyset \\
 \text{reorderer_state}' = \text{ReordererFinishedState} \\
 \text{---}
 \end{array}$$

StepReorderer steps the reorderer.

$$\begin{array}{l}
 \text{StepReorderer} \hat{=} \text{StartReorderer} \vee \\
 \quad \text{ReordererAddBestCondition} \vee \text{ReordererAddBestTiedCondition} \vee \\
 \quad \text{FinishReorderer}
 \end{array}$$

ResetReorderer allows the system to reset the reorderer to the finished state from any state.

$$\begin{array}{l}
 \text{--- } \underline{\text{ResetReorderer}} \text{ ---} \\
 \Delta \text{Reorderer} \\
 \text{---} \\
 \text{reorderer_state}' = \text{ReordererFinishedState} \\
 \text{---}
 \end{array}$$

For simplicity, the reordering of conjunctive negations has not been specified. Whenever a conjunctive negation is added to the ordering the system should reorder the conditions of the negation using the same algorithm. It should start with: an empty ordering, the bound variables of the entire ordering, the roots of the entire ordering and the positively tested variables of the positive conditions of the conjunctive negation.

Bibliography

- [Aky90] A. Akyurek. Means-Ends Planning, Operator Subgoalng, and Operator Valuation: An Example Soar Program. Technical Report RUG-FA-90-3, University of Groningen, August 1990.
- [BW90] Michael Barr and Charles Wells. *Category Theory for Computing Science*. Prentice Hall International, 66 Woodlane End, Hemmel Hampstead, Hertfordshire HP24RG UK, 1990.
- [Cra91] Iain D Craig. *Formal Specification of Advanced AI Architectures*. Ellis Horwood Limited, Market Cross House, Cooper Street, Chichester, West Sussex, PO19 1EB, England, 1991.
- [DG89] N Delisle and D Garlan. Formally specifying electronic instruments. In *Proceedings of Fifth International Workshop on Software Specification and Design*, pages 242-248, 19-20 May 1989.
- [DS90] M. DeJongh and Jr. Smith, J. W. Blood Typing: Functional Modeling Put to the Tests. Laboratory for Artificial Intelligence Research, Ohio State University, November, 1990, Unpublished., 1990.
- [Gar91] David Garlan. Formalizing design spaces: implicit invocation mechanisms. Technical Report CMU-CS-91-114, Carnegie Mellon University, School of Computer Science, 1991.
- [Gut90] John V. Guttag. Report on the Larch Shared Language. Technical Report 58, Digital Equipment Corporation, Systems Research Center, 1990.
- [GW88] Joseph A Goguen and Timothy Winkler. Introducing OBJ3. Technical Report SRI-CSL-88-9, SRI International, Computer Science Laboratory, 1988.
- [Hay85] Ian Hayes, editor. *Specification Case Studies*. Prentice Hall International, 66 Woodlane End, Hemmel Hampstead, Hertfordshire HP24RG UK, 1985.

- [HPS89] W. Hsu, M. Prietula, and D.M. Steier. Merl-Soar: Scheduling within a general architecture for intelligence. In *Proceedings of the Third International Conference on Expert Systems and the Leading Edge in Production and Operations Management*, pages 467-481, May 1989.
- [Jon90] C.B. Jones. *Systematic Software Development using VDM (2nd edition)*. Prentice Hall International, 66 Woodlane End, Hemmel Hampstead, Hertfordshire HP24RG UK, 1990.
- [JRS91] B.E. John, R.W. Remington, and D.M. Steier. An Analysis of Space Shuttle Countdown Activities: Preliminaries to a Computational Model of the NASA Test Director. Technical Report CMU-CS-91-138, School of Computer Science, Carnegie Mellon University, May 1991.
- [JS90] C. B. Jones and R. (editors) Shaw. *Case Studies in VDM*. Prentice Hall International, 66 Woodlane End, Hemmel Hampstead, Hertfordshire HP24RG UK, 1990.
- [KR88] Brian W. Kernighan and Dennis M. Richie. *The C Programming Language*. Prentice Hall, Englewood Cliffs, New Jersey 07632, 1988.
- [Lai86] J.E. Laird. Soar User's Manual: Version 4.0. Technical report, Intelligent Systems Laboratory, Palo Alto Research Center, Xerox Corporation, January 1986. Out of Print, see the Version 5.2 manual.
- [LCAS90] J.E. Laird, C.B. Congdon, E. Altmann, and K. Swedlow. Soar User's Manual: Version 5.2. Technical report, Electrical Engineering and Computer Science, University of Michigan, October 1990. Also available from The Soar Group, School of Computer Science, Carnegie Mellon University, CMU-CS-90-179.
- [LLN91] J. Fain Lehman, R.L. Lewis, and A. Newell. Integrating Knowledge Sources in Language Comprehension. School of Computer Science, Carnegie Mellon University, January 1991, Unpublished, 1991.
- [LNR87] J.E. Laird, A. Newell, and P.S. Rosenbloom. Soar: An architecture for general intelligence. *Artificial Intelligence*, 33(1):1-64, 1987.
- [LRN84] J.E. Laird, P.S. Rosenbloom, and A. Newell. Towards chunking as a general learning mechanism. In *Proceedings of the National Conference on Artificial Intelligence*, pages 188-192, August 1984.
- [LRN86] J. E. Laird, P.S. Rosenbloom, and A. Newell. Chunking in Soar: The anatomy of a general learning mechanism. *Machine Learning*, 1(1):11-46, 1986.

- [LYHT91] J.E. Laird, E.S. Yager, M. Hucka, and C.M. Tuck. Robo-Soar: An integration of external interaction, planning, and learning using Soar. In Van de Velde, editor, *Machine Learning for Autonomous Agents*. MIT Press: Bradford Books, Cambridge, Massachusetts, 1991.
- [New90] A. Newell. *Unified Theories of Cognition*. Harvard University Press, Cambridge, Massachusetts, 1990.
- [NYL+91] A. Newell, G.R. Yost, Laird, P.S. J.E., Rosenbloom, and E. Altmann. Formulating the problem space computational model. In R.F. Rashid, editor, *Carnegie Mellon Computer Science. A 25-Year Commemorative*, pages 255-293. ACM-Press: Addison-Wesley, Reading, PA, 1991.
- [PMN91] G. Pelton, B. G. Milnes, and A. Newell. NSoar External Interface Informal Specification, 1991. NSoar Group, School of Computer Science, Carnegie Mellon University, May, 1991, Unpublished.
- [PTS91] Ben Potter, David Till, and Jane Sinclair. *An Introduction to formal specification and Z*. Prentice Hall International, 66 Woodlane End, Hemmel Hampstead. Hertfordshire HP24RG UK, 1991
- [RL86] P. S. Rosenbloom and J. E. Laird. Mapping explanation-based generalization onto soar. In *Proceedings of the National Conference on Artificial Intelligence*, pages 561-567, August 1986.
- [Sca86] D.J. Scales. Efficient matching algorithms for the SOAR/OPS5 production system. Technical Report KSL-86-47, Knowledge Systems Laboratory, Stanford University, June 1986.
- [SK91] A. Simon, T. and Newell and D. Klahr. Q-Soar: A computational account of children's learning about number conservation. In *Working Models of Human Perception*. Morgan Kaufman, Los Altos, California, 1991.
- [Spi88] J.M. Spivey. *Understanding Z: a specification language and its formal semantics*. Cambridge University Press, Cambridge, Cambridgeshire, UK, 1988.
- [Spi89] J.M. Spivey. *The Z Notation*. Prentice Hall, Hertfordshire, UK, 1989.
- [Tam91] M. Tambe. *Eliminating Combinatorics from Production Match*. PhD thesis, School of Computer Science, Carnegie Mellon University, May 1991.
- [Wor92] J. B. Wordsworth. *Software Development with Z* International Computer Science series. Addison Wesley, 1992. ISBN 0-210-62757-4.

List of Definitions

1	Introduction	10
2	Specifying Soar's Control: State Machines	18
1	<i>StateSchema</i>	19
2	<i>SMInitialState</i>	19
3	<i>SMS1State</i>	19
4	<i>SMS2State</i>	19
5	<i>SMS3State</i>	19
6	<i>SMFinishedState</i>	19
7	<i>StateCounter</i>	21
8	<i>SMInitialState</i>	21
9	<i>SMS1State</i>	21
10	<i>SMS2State</i>	21
11	<i>SMS3State</i>	21
12	<i>SMFinishedState</i>	21
13	<i>StateSchema</i>	22
14	<i>InitStateSchema</i>	22
15	<i>InitStateSchema</i>	23
16	<i>SMOperation</i>	23
17	<i>SMOperation</i>	24
18	<i>SMOperation</i>	24
19	<i>SMInitialize</i>	24
20	<i>SMStartS1</i>	25
21	<i>SMS1</i>	25
22	<i>SMS2OP1</i>	26
23	<i>SMS2OP2</i>	26
24	<i>SMS2</i>	26
25	<i>SMStartS3</i>	27
26	<i>SMStepS3</i>	27
27	<i>SMS3</i>	27

28	<i>SMFinish</i>	28
29	<i>SMStep</i>	28
30	<i>SMReset</i>	28

3 Base Symbol Structures 30

31	Σ	30
32	<i>Variable</i>	30
33	<i>Identifier</i>	30
34	<i>Constant</i>	30
35	<i>VariableOrConstant</i>	30
36	<i>Symbol</i>	30
37	'@'	31
38	'!'	31
39	'+'	31
40	'-'	31
41	'.'	31
42	'>'	31
43	'<'	31
44	'=P'	31
45	'&'	31
46	'=R'	31
47	'<>'	31
48	'<=>'	31
49	'('	31
50	')'	31
51	'	31
52	'-- >'	31
53	<i>PreferenceSymbol</i>	31
54	<i>RelationSymbol</i>	31
55	<i>UnaryPreferenceSymbol</i>	31
56	<i>BinaryPreferenceSymbol</i>	31
57	'NIL'	31
58	'GOAL'	31
59	'PROBLEM-SPACE'	31
60	'STATE'	31
61	'OPERATOR'	31
62	'IMPASSE'	31
63	'OBJECT'	31
64	'ITEM'	31
65	'ATTRIBUTE'	31
66	'CHOICES'	31
67	'CONSTRAINT-FAILURE'	31
68	'CONFLICT'	31

69	'TIE'	31
70	'NO-CHANGE'	31
71	'NONE'	31
72	'MULTIPLE'	31
73	'QUIESCENCE'	31
74	'TYPE'	31
4	Temporary Memory and Production Memory	33
75	<i>UnaryPreference</i>	33
76	<i>UnaryPreferencesComponents</i>	34
77	<i>BinaryPreference</i>	34
78	<i>BinaryPreferencesComponents</i>	34
79	<i>Preference</i>	34
80	<i>UnaryP</i>	34
81	<i>BinaryP</i>	34
82	<i>PreferencesComponents</i>	35
83	<i>makeUnaryPreference</i>	35
84	<i>makeBinaryPreference</i>	35
85	<i>PreferencesId</i>	36
86	<i>ConnectedPreference</i>	36
87	<i>TCIPreference</i>	37
88	<i>OfIdPreference</i>	37
89	<i>YesOrNo</i>	37
90	<i>Yes</i>	37
91	<i>No</i>	37
92	<i>WME</i>	37
93	<i>WMEsComponents</i>	38
94	<i>makeWME</i>	38
95	<i>ConnectedWME</i>	38
96	<i>TCIWME</i>	38
97	<i>OfIdWME</i>	38
98	<i>TME</i>	39
99	<i>PreferenceTME</i>	39
100	<i>WMETME</i>	39
101	<i>TMEsComponents</i>	39
102	<i>ConnectedTME</i>	39
103	<i>TCITME</i>	39
104	<i>OfIdTME</i>	40
105	<i>UnaryMake</i>	40
106	<i>UnaryMakesComponents</i>	40
107	<i>BinaryMake</i>	40
108	<i>BinaryMakesComponents</i>	41
109	<i>Make</i>	41

110	<i>UnaryM</i>	41
111	<i>BinaryM</i>	41
112	<i>MakesComponents</i>	41
113	<i>makeUnaryMake</i>	42
114	<i>makeBinaryMake</i>	42
115	<i>ConnectedMake</i>	42
116	<i>TCIMake</i>	43
117	<i>OfIdMake</i>	43
118	<i>Test</i>	43
119	<i>BlankTest</i>	43
120	<i>YesTest</i>	43
121	<i>NoTest</i>	43
122	<i>EqualityTest</i>	43
123	<i>NotTest</i>	43
124	<i>SameTypeTest</i>	43
125	<i>DisjunctiveTest</i>	43
126	<i>ConjunctiveTest</i>	43
127	<i>TestsComponents</i>	45
128	<i>TestsPositiveComponents</i>	45
129	<i>SimpleTest</i>	46
130	<i>WMETest</i>	46
131	<i>makeWMETest</i>	46
132	<i>WMETestsComponents</i>	46
133	<i>WMETestsPositiveComponents</i>	47
134	<i>ConnectedWMETest</i>	47
135	<i>Condition</i>	47
136	<i>PositiveCondition</i>	47
137	<i>NegativeCondition</i>	47
138	<i>NegativeConjunctiveCondition</i>	47
139	<i>SimpleCondition</i>	48
140	<i>ConditionsComponents</i>	48
141	<i>ConditionsPositiveComponents</i>	48
142	<i>TCICondition</i>	50
143	<i>OfIdCondition</i>	50
144	<i>TestedVariablesBound</i>	51
145	<i>SP</i>	52
146	<i>makeSP</i>	53
147	<i>SpsComponents</i>	53

5 Recognition Memory 54

148	<i>Binding</i>	55
149	<i>Consistent</i>	55
150	<i>ConsistentExtension</i>	55

151	<i>Covers</i>	55
152	<i>Matches</i> _{Test}	56
	<i>Symbol</i>	
153	<i>Matches</i> _{WME} <i>Test</i>	57
154	<i>Match</i>	58
155	<i>makeMatch</i>	58
156	<i>Matches</i> _{Condition}	59
	<i>WM</i>	
157	<i>Matches</i> _{Match}	60
	<i>WM</i>	
158	<i>Instantiate</i> <i>VariableOrConstant</i>	60
159	<i>Instantiate</i> <i>SimpleTest</i>	61
160	<i>Instantiate</i> <i>SetOfSimpleTests</i>	61
161	<i>Instantiate</i> <i>Test</i>	62
162	<i>Instantiate</i> <i>WME</i> <i>Test</i>	62
163	<i>Instantiate</i> <i>Condition</i>	63
164	<i>Instantiate</i> <i>SetOfConditions</i>	63
165	<i>Instantiate</i> <i>Conditions</i> <i>WME</i>	64
166	<i>Instantiate</i> <i>Condition</i> <i>Matching</i>	65
167	<i>Makes</i> <i>Preference</i>	66
168	<i>Instantiation</i>	67
169	<i>make</i> <i>Instantiation</i>	67
170	<i>Goal</i> <i>State</i>	68
171	<i>Goal</i> <i>Operator</i>	68
172	<i>Goal</i> <i>Ancestors</i>	68
173	<i>Match</i> <i>Goals</i>	69
174	<i>Match</i> <i>Goal</i>	69
175	<i>Not</i>	70
176	<i>Tests</i> <i>Not</i>	71
177	<i>WME</i> <i>Tests</i> <i>Not</i>	71
178	<i>Conditions</i> <i>Not</i>	71
179	<i>Assignment</i>	72
180	<i>Variabilize</i> <i>Symbol</i>	72
181	<i>Variabilize</i> <i>Test</i>	72
182	<i>Variabilize</i> <i>WME</i> <i>Test</i>	73
183	<i>Variabilize</i> <i>Condition</i>	73
184	<i>Variabilize</i> <i>Condition</i> <i>Matching</i>	74
185	<i>Variabilize</i> <i>Preference</i> <i>ToMake</i>	75
186	<i>Variabilize</i> <i>RHS</i>	75
187	<i>VNot</i>	76
188	<i>Not</i> <i>ToVNot</i>	76
189	<i>first</i> ₃	76
190	<i>second</i> ₃	76
191	<i>third</i> ₃	76
192	<i>Notify</i> <i>Test</i>	77
193	<i>Notify</i> <i>WME</i> <i>Test</i>	78

194	<i>NotifyCondition</i>	78
195	<i>NotifyConditionMatchings</i>	79
196	<i>Grounds</i>	80
197	<i>Potentials</i>	80
198	<i>Locals</i>	80
199	<i>Nots</i>	80
200	<i>Trace</i>	81
201	<i>makeTrace</i>	81
202	<i>TracesComponents</i>	81
203	<i>ChunkingState</i>	83
204	<i>ChunkingInitialState</i>	83
205	<i>ChunkingTraceConditionState</i>	83
206	<i>ChunkingTraceUnGroundedPotentialsState</i>	83
207	<i>ChunkingFinishedState</i>	83
208	<i>Chunking</i>	84
209	Δ <i>Chunking</i>	86
210	<i>InitChunking</i>	87
211	<i>ChunkingInitialize</i>	87
212	<i>ChunkingReset</i>	88
213	<i>Results</i>	88
214	<i>ChunkingStartTracing</i>	89
215	<i>ChunkingIgnoreInstantiationInTopGoal</i>	89
216	<i>ChunkingNoResults</i>	90
217	<i>Owner</i>	90
218	<i>TraceWME</i>	91
219	<i>TraceItems</i>	91
220	<i>ChunkingTraceSeed</i>	92
221	<i>ChunkingSkipSeed</i>	93
222	<i>ChunkingTraceLocal</i>	93
223	<i>ImpasseAttribute</i>	94
224	<i>ChunkingSkipArchitectureGeneratedLocal</i>	94
225	<i>ChunkingTraceQuiescenceLocal</i>	95
226	<i>ChunkingPotentializeLocal</i>	96
227	<i>ChunkingTraceGroundedPotential</i>	96
228	<i>ChunkingStartTraceUnGroundedPotentials</i>	97
229	<i>ChunkingTraceUnGroundedPotential</i>	98
230	<i>ChunkingSkipUnGroundedPotential</i>	98
231	<i>ChunkingTraceUnGroundedPotentialsDone</i>	99
232	<i>ChunkingFinishLearnOn</i>	100
233	<i>ChunkingFinishLearnOff</i>	101
234	<i>ChunkingStep</i>	101
235	<i>Slot</i>	101
236	<i>SlotsWMEs</i>	102
237	<i>RM</i>	105

238	<i>InitRM</i>	106
239	<i>InstantiationsLHSOperators</i>	107
240	<i>InstantiationsRHSOperators</i>	107
241	<i>OperatorCreationLHS</i>	107
242	<i>OperatorCreationRHS</i>	108
243	<i>OperatorCreation</i>	108
244	<i>OperatorModificationLHS</i>	108
245	<i>OperatorModificationRHS</i>	109
246	<i>OperatorModification</i>	109
247	<i>OperatorApplicationLHS</i>	109
248	<i>OperatorApplicationRHS</i>	110
249	<i>OperatorApplication</i>	110
250	<i>OSupport</i>	110
251	<i>PPhaseInitialize</i>	111
252	<i>PPhaseReset</i>	111
253	<i>Quiescence</i>	111
254	<i>PPhaseStartRetractInstantiation</i>	111
255	<i>PPhaseRetractInstantiation</i>	112
256	<i>PPhaseStartInstantiateMatch</i>	112
257	<i>PPhaseInstantiateMatch</i>	113
258	<i>PPhaseConferSupport</i>	114
259	<i>PPhaseStepChunking</i>	114
260	<i>PPhaseChunkingReturnedInstantiation</i>	115
261	<i>PPhaseChunkingNoInstantiation</i>	115
262	<i>PPhaseStartOReject</i>	115
263	<i>PPhaseORejectPreferences</i>	116
264	<i>PPhaseStartChangePM</i>	116
265	<i>PPhaseAddPreference</i>	117
266	<i>PPhaseRemovePreference</i>	117
267	<i>PPhaseFinish</i>	117

6 IO 119

268	<i>InputChannel</i>	119
269	<i>InputAttribute</i>	119
270	<i>OutputChannel</i>	119
271	<i>OutputAttribute</i>	119
272	<i>InputMapping</i>	121
273	<i>OutputMapping</i>	121
274	<i>OutputStructure</i>	121
275	<i>NewInputStructure</i>	121
276	<i>ModifyInputStructure</i>	122
277	<i>DeleteInputStructure</i>	122
278	<i>InputStructure</i>	122

279	<i>New</i>	122
280	<i>Modify</i>	122
281	<i>Delete</i>	122
282	<i>CycleOfInput</i>	122
283	<i>CycleOfOutput</i>	122
284	<i>InputCycleState</i>	123
285	<i>InputCycleInitialState</i>	123
286	<i>InputCycleReadState</i>	123
287	<i>InputCycleFinishedState</i>	123
288	<i>OutputCycleState</i>	124
289	<i>OutputCycleInitialState</i>	124
290	<i>OutputCycleSendState</i>	124
291	<i>OutputCycleFinishedState</i>	124
292	<i>IO</i>	125
293	<i>InitIO</i>	125
294	<i>InputLink</i>	126
295	<i>LegalInput</i>	127
296	<i>InputCycleInitialize</i>	128
297	<i>StartReadInput</i>	128
298	<i>ReadNewInputStructureObject</i>	128
299	<i>ReadNewInputStructureWME</i>	129
300	<i>ReadNewInputStructureDone</i>	129
301	<i>ReadNewInputStructure</i>	130
302	<i>ReadModifyInputStructureAdd</i>	130
303	<i>ReadModifyInputStructureDelete</i>	131
304	<i>ReadModifyInputStructureDone</i>	131
305	<i>ReadModifyInputStructure</i>	131
306	<i>ReadDeleteInputStructure</i>	132
307	<i>CloseInputChannel</i>	132
308	<i>ReadInputChannel</i>	133
309	<i>InputCycleFinish</i>	133
310	<i>InputCycleStep</i>	133
311	<i>InputCycleReset</i>	133
312	<i>OutputCycleInitialize</i>	134
313	<i>StartSendOutputChannel</i>	134
314	<i>SendOutputChannelOutputStructureObject</i>	135
315	<i>SendOutputChannelOutputStructureWME</i>	135
316	<i>SendOutputChannel</i>	135
317	<i>OutputCycleFinish</i>	136
318	<i>OutputCycleStep</i>	136
319	<i>OutputCycleReset</i>	136

7	Decide	137
320	<i>Impasses WMEs</i>	138
321	<i>GoalDescendants</i>	139
322	<i>ImpasserState</i>	139
323	<i>ImpasserInitialState</i>	139
324	<i>ImpasserCreateOrChangeState</i>	139
325	<i>ImpasserFinishedState</i>	139
326	<i>ImpasseType</i>	141
327	<i>ContextChanged</i>	141
328	<i>Changed</i>	141
329	<i>Unchanged</i>	141
330	<i>ImpasserStateSchema</i>	141
331	<i>PreferenceSemanticsState</i>	142
332	<i>PSInitialState</i>	142
333	<i>PSRequireState</i>	142
334	<i>PSAcceptableState</i>	142
335	<i>PSProhibitState</i>	142
336	<i>PSRejectState</i>	142
337	<i>PSBetterWorseState</i>	142
338	<i>PSBestState</i>	142
339	<i>PSWorstState</i>	142
340	<i>PSIndifferentState</i>	142
341	<i>PSParallelState</i>	142
342	<i>PSFinishedState</i>	142
343	<i>NumberOfWinners</i>	144
344	<i>All</i>	144
345	<i>One</i>	144
346	<i>PreferenceSemanticsStateSchema</i>	144
347	<i>WMPhaseState</i>	146
348	<i>WMPhaseInitialState</i>	146
349	<i>WMPhasePickSlotState</i>	146
350	<i>WMPhaseDecideSlotState</i>	146
351	<i>WMPhaseImpasserState</i>	146
352	<i>WMPhaseChangeWMState</i>	146
353	<i>WMPhaseFinishedState</i>	146
354	<i>WMPhaseStateSchema</i>	146
355	<i>QPhaseState</i>	148
356	<i>QPhaseInitialState</i>	148
357	<i>QPhasePickSlotState</i>	148
358	<i>QPhaseDecideSlotState</i>	148
359	<i>QPhaseImpasserState</i>	148
360	<i>QPhaseChangeWMState</i>	148
361	<i>QPhaseFinishedState</i>	148

362	<i>ContextSlotAttribute</i>	148
363	<i>QPhaseStateSchema</i>	148
364	<i>ContextImpassePreferences</i>	149
365	<i>Decide</i>	150
366	<i>InitDecide</i>	150
367	<i>DecideCreateFirstGoal</i>	151
368	<i>ImpasserInitialize</i>	152
369	<i>ImpasserRemoveImpasse</i>	152
370	<i>ImpasserNoImpasseToRemove</i>	153
371	<i>ImpasserCreateImpasse</i>	154
372	<i>ImpasserChangeImpasse</i>	155
373	<i>ImpasserNoNewImpasseOrChanges</i>	155
374	<i>StepImpasser</i>	156
375	<i>PSInitialize</i>	156
376	<i>PSStart</i>	156
377	<i>PSOneRequire</i>	157
378	<i>PSMultipleRequires</i>	157
379	<i>PSRequireProhibited</i>	158
380	<i>PSNoRequires</i>	158
381	<i>PSRequire</i>	158
382	<i>PSAcceptable</i>	159
383	<i>PSNoAcceptable</i>	159
384	<i>PSProhibit</i>	160
385	<i>PSReject</i>	160
386	<i>Conflict</i>	161
387	<i>PSNoBetterWorseConflict</i>	161
388	<i>PSBetterWorseConflict</i>	162
389	<i>PSBetterWorse</i>	162
390	<i>PSBest</i>	162
391	<i>PSWorst</i>	163
392	<i>AllMutuallyIndifferent</i>	163
393	<i>PSMutuallyIndifferent</i>	164
394	<i>PSNotMutuallyIndifferentNonContextSlot</i>	164
395	<i>PSNotMutuallyIndifferentContextSlot</i>	165
396	<i>PSIndifferent</i>	165
397	<i>AllMutuallyParallel</i>	165
398	<i>PSMutuallyParallel</i>	166
399	<i>PSNotMutuallyParallel</i>	166
400	<i>PSParallel</i>	166
401	<i>PSStep</i>	166
402	<i>WMPhaseInitialize</i>	167
403	<i>WMPhaseReset</i>	168
404	<i>WMPhaseStartPickSlot</i>	168
405	<i>WMPhaseStartDecideSlot</i>	168

406	<i>WMPhaseStepDecideSlot</i>	169
407	<i>WMPhaseStartImpasser</i>	169
408	<i>WMPhaseStepImpasser</i>	169
409	<i>WMPhaseImpasserFinished</i>	170
410	<i>WMPhaseAddWME</i>	171
411	<i>WMPhaseRemoveWME</i>	172
412	<i>WMPhaseChangeWMFinished</i>	172
413	<i>WMPhaseFinish</i>	173
414	<i>WMPhaseStep</i>	173
415	<i>QPhaseInitialize</i>	174
416	<i>QPhaseReset</i>	174
417	<i>QPhaseStartPickSlot</i>	175
418	<i>NextContextSlot</i>	175
419	<i>LastSlotToCheckInContext</i>	175
420	<i>LastSlotToCheckInContextStack</i>	176
421	<i>DecidableSlot</i>	176
422	<i>QPhaseNextSlot</i>	177
423	<i>QPhaseNoChangedPreferencesNoChangeImpasse</i>	178
424	<i>QPhaseStartDecideSlot</i>	178
425	<i>QPhaseStepDecideSlot</i>	179
426	<i>QPhaseStartImpasser</i>	179
427	<i>QPhaseChangedPreferencesNoChangeImpasse</i>	180
428	<i>QPhaseStepImpasser</i>	180
429	<i>QPhaseImpasserFinished</i>	180
430	<i>LaterSlotsOfContextSlot</i>	181
431	<i>QPhaseAddWME</i>	181
432	<i>QPhaseRemoveWME</i>	182
433	<i>QPhaseReAddWME</i>	183
434	<i>QPhaseSlotUnchanged</i>	183
435	<i>QPhaseSlotChanged</i>	184
436	<i>QPhaseStep</i>	184
8	Top Level	185
437	<i>TopLevelState</i>	186
438	<i>TLInitialState</i>	186
439	<i>TLDecisionCycleState</i>	186
440	<i>TLElaborationPhaseState</i>	186
441	<i>TLElaborationCycleState</i>	186
442	<i>TLInputCycleState</i>	186
443	<i>TLPreferencePhaseState</i>	186
444	<i>TLWorkingMemoryPhaseState</i>	186
445	<i>TLOutputCycleState</i>	186
446	<i>TLDecisionPhaseState</i>	186

447	<i>TLFinishedState</i>	186
448	<i>TL</i>	186
449	<i>Soar</i>	188
450	<i>InitSoar</i>	188
451	<i>TLInitialize</i>	188
452	<i>TLReset</i>	189
453	<i>TLStartSoar</i>	189
454	<i>TLStartDecisionCycle</i>	189
455	<i>TLStartElaborationPhase</i>	189
456	<i>TLStartElaborationCycle</i>	190
457	<i>TLStepInputCycle</i>	190
458	<i>TLStartPreferencePhase</i>	190
459	<i>TLStartDecisionPhase</i>	191
460	<i>TLStepPreferencePhase</i>	191
461	<i>TLStartWorkingMemoryPhase</i>	191
462	<i>TLStepWorkingMemoryPhase</i>	191
463	<i>TLStartOutputCycle</i>	192
464	<i>TLStepOutputCycle</i>	192
465	<i>TLEndElaborationCycle</i>	192
466	<i>TLStepQPhase</i>	192
467	<i>TLEndDecisionCycle</i>	193
468	<i>TLStep</i>	193
A Numbers		194
469	<i>Number</i>	194
470	<i>></i>	194
471	<i><</i>	194
472	<i>>=</i>	194
473	<i><=</i>	194
474	<i>NumericRelationSymbols</i>	194
475	<i>RelationalTest</i>	194
476	<i>Matches^{Test}_{Symbol}</i>	194
B An Implementation Discipline		196
C Finer Grained Hooks		201
477	<i>AddWME</i>	202
478	<i>RemoveWME</i>	202
479	<i>AddPreference</i>	202
480	<i>RemovePreference</i>	202
481	<i>AddTME</i>	203

482	<i>RemoveTME</i>	203
483	<i>AddSP</i>	203
484	<i>RemoveSP</i>	204
485	<i>RemoveContextImpasse</i>	204
486	<i>RemoveNonContextImpasse</i>	205

D The Reorderer 206

487	<i>C</i>	207
488	<i>B</i>	207
489	<i>U</i>	207
490	<i>GIU</i>	207
491	<i>GIB</i>	207
492	<i>YES</i>	207
493	<i>NO</i>	207
494	<i>TestClass</i>	207
495	<i>ClassifySigma</i>	207
496	<i>ClassifyTest</i>	208
497	<i>WMETestClass</i>	209
498	<i>WC</i>	209
499	<i>ClassifyWMETest</i>	209
500	<i>G</i>	210
501	<i>O</i>	210
502	<i>A</i>	210
503	<i>V</i>	210
504	<i>AP</i>	210
505	<i>ABF</i>	214
506	<i>ABFI</i>	214
507	<i>ABFA</i>	214
508	<i>ABFVAP</i>	214
509	<i>UBF</i>	215
510	<i>DBF</i>	215
511	<i>BF</i>	216
512	<i>ContainsPositiveConstants</i>	216
513	<i>SBF</i>	217
514	<i>ConditionsBranchingFactor</i>	218
515	<i>Reorderer</i>	220
516	<i>InitReorderer</i>	220
517	Δ <i>Reorderer</i>	221
518	<i>InitializeReorderer</i>	221
519	<i>StartReorderer</i>	221
520	<i>BestConditions</i>	222
521	<i>ReordererAddBestCondition</i>	222
522	<i>LookAhead</i>	223

523	<i>BestTiedConditions</i>	223
524	<i>ReordererAddBestTiedCondition</i>	224
525	<i>FinishReorderer</i>	224
526	<i>ResetReorderer</i>	224