

Synthesis of High-Performance Transportation Schedulers

Douglas R. Smith, Eduardo A. Parra, Stephen J. Westfold
Kestrel Institute
3260 Hillview Avenue
Palo Alto, California 94304
9 March 1995

Abstract

This report describes our research on transportation planning and scheduling supported by the ARPA/Rome Lab Planning Initiative (ARPI). The main goal of this project was to develop generic tools to support the construction of flexible, high-performance planning and scheduling software. Our technical approach is based on *program transformation* technology which allows the systematic machine-supported development of software from requirement specifications. The development process can produce highly efficient code along with a proof of the code's correctness.

We have used KIDS (Kestrel Interactive Development System) to derive extremely fast and accurate transportation schedulers from formal specifications. As test data we use strategic transportation plans which are generated by U.S. government planners. A typical problem, with 10,000 movement requirements, takes the derived scheduler 1 – 3 minutes to solve, compared with 2.5 *hours* for a deployed feasibility estimator (JFAST) and 36 *hours* for deployed schedulers (FLOGEN, ADANS). The computed schedules use relatively few resources and satisfy all specified constraints. The speed of this scheduler is due to the synthesis of strong constraint checking and constraint propagation code.

Contents

1	Introduction	4
2	Overview of Results	6
3	KIDS model of program development	8
4	Specifying a Scheduler	9
4.1	What is Scheduling?	9
4.2	Strategic Transportation Scheduling	10
4.3	(Re-)Formulating Domain Theories for Transportation Scheduling	11
4.4	Formal Specification of a Scheduler	12
5	Synthesizing a Scheduler	14
5.1	Approach	14
5.1.1	Problem Theories	14
5.1.2	Algorithm Theories	16
5.2	Synthesizing a Scheduler	16
5.2.1	Global Search Theory	16
5.2.2	Pruning Mechanisms	19
5.2.3	Cutting Constraints and Constraint Propagation	20
5.2.4	Constraint Relaxation	28
5.2.5	Using KIDS	28
5.3	KTS – Strategic Transportation Scheduling	29
5.4	ITAS – In-Theater Airlift Scheduler	30

6	Classification Approach to Algorithm Design	31
6.1	Technical Foundations – Theories	31
6.2	Refinement Hierarchy and the Ladder Construction	31
6.3	Constructing Theory Morphisms	34
7	Concluding Remarks	36
	References	36
A	Relation-based Domain Theory for Transportation Scheduling	41
B	Map-based Domain Theory for Transportation Scheduling	46
C	Global Search Theory for Transportation Scheduling	68
D	Derived Transportation Scheduler	70

1 Introduction

This report describes our research on the transformational development of high-performance transportation schedulers. Our approach to developing scheduling software involves several stages. The first step is to develop a formal model of the transportation scheduling domain, called a *domain theory*. Second, the constraints, objectives, and preferences of a particular scheduling problem are stated within a domain theory as a *problem specification*. Finally, an executable scheduler is produced semi-automatically by applying a sequence of *transformations* to the problem specification. The transformations embody programming knowledge about algorithms, data structures, program optimization techniques, etc. The result of the transformation process is executable code that is consistent with the given problem specification. Furthermore, the resulting code can be extremely efficient.

Transportation scheduling tools currently used by the U.S. government are based on models of the transportation domain that few people understand [11]. Consequently, users often do not trust that the scheduling results reflect the characteristics of the current situation. Our approach tries to address this issue by making the domain model and scheduling problem explicit and clear. If a scheduling situation arises which is not treated by existing scheduling tools, the user can specify the problem and generate a situation-specific scheduler.

One of the benefits of a transformational approach to scheduling is the synthesis of specialized constraint management code. Previous systems for performing scheduling in AI (e.g. [13, 12, 48, 47]) and Operations Research [2, 24] use constraint representations and operations that are geared for a broad class of problems, such as constraint satisfaction problems or linear programs. In contrast, transformational techniques can derive specialized representations for constraints and related data, and also derive efficient specialized code for constraint checking and constraint propagation.

Figure 1 describes our vision of an advanced environment for producing planning/scheduling software. Briefly, the idea is to rapidly develop a situation-specific domain model and problem specification using a knowledge-elicitation system, and then to *synthesize* high-performance planning and scheduling tools that are specialized to the current situation. The majority of users' interaction would be codifying the *domain theory* and specification of the current situation, to aid in synthesizing a customized planning/scheduling tool.

We now step through the process in more detail. Several classes of users are involved in the construction and use of a scheduling system.

One class of users, who include domain experts and specialists in model construction, interact with a knowledge elicitation system to help classify the features of the situation and select, compose, extend, and refine, (possibly abstract) models from a preexisting library of domain models. The result is a model and problem specification tailored to the details of the current situation (as closely as expertise and time permit).

Another class of users, who specialize in software design and formal modeling of programming knowledge, interact with a planning/scheduling synthesis system to develop code from the problem specification. The interaction involves composing components from a library of

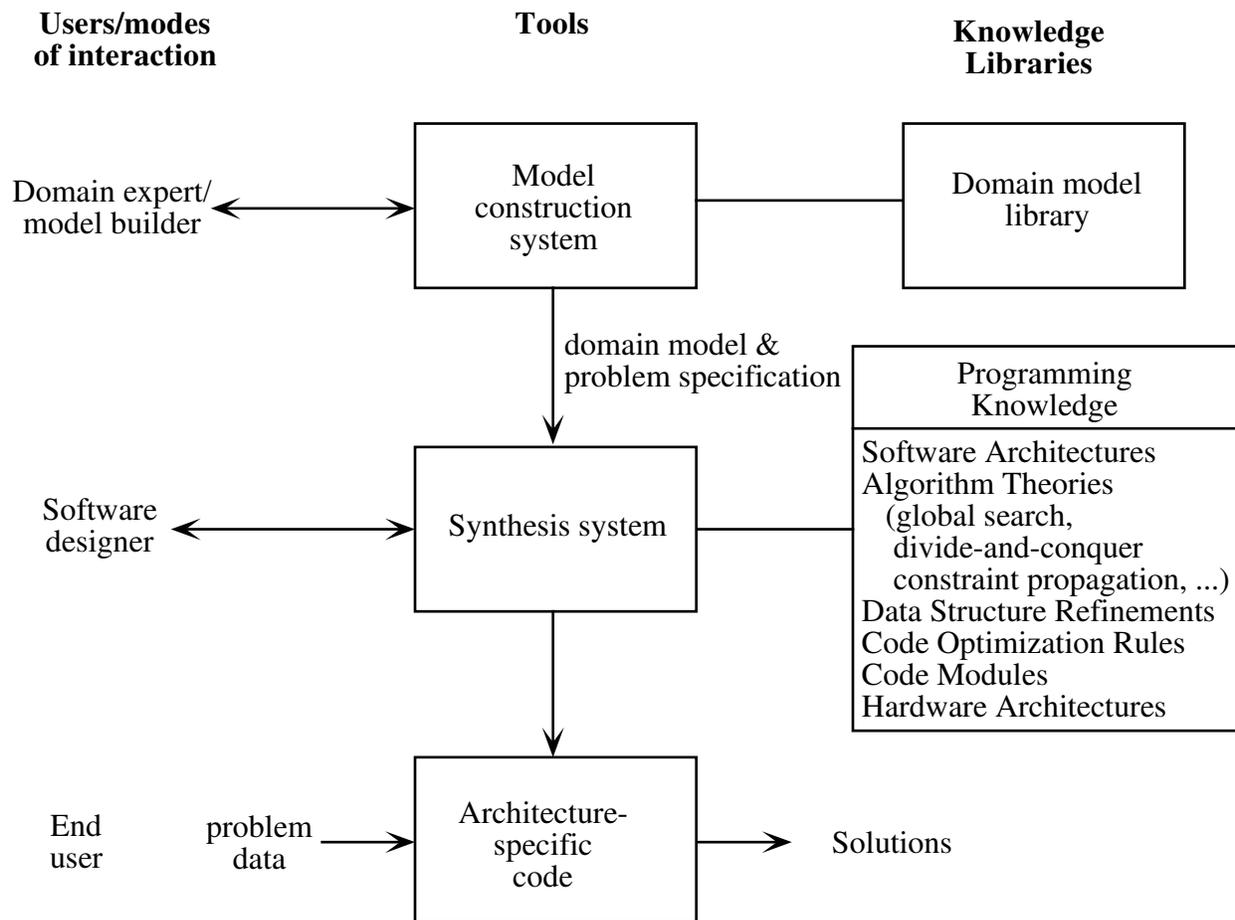


Figure 1: Advanced Development Environment for Planning and Scheduling Software

reusable parts, or selecting and applying representations of abstract programming knowledge about algorithms, data structures, code optimization techniques, planning and scheduling-specific design strategies, and so on. Most of the design process is automated, with only a few high-level design decisions made by the developer. Another interactive task for this user is the evolution of the programming knowledge-base itself.

The output of the synthesis system is executable planning/scheduling code which can then be embedded in a planning/scheduling system and executed by an end-user.

There is a feedback loop implicit in the figure. The end user/domain expert using the synthesized code may detect missing constraints, or have preferences or other information not accounted for in the code. This information is fed back to the model-building stage and the process iterates. The fact that each synthesis step preserves consistency between problem specification and generated code means that *maintenance and evolution back up to the specification/domain model level*, not the code-level as in current practice.

We developed an approximation to this vision in the current project, based on the KIDS system, and demonstrated its feasibility.

In Sections 3 through 5 we present KIDS and the process of developing domain theories, specifications, and code for scheduling problems. Section 7 describes fundamental work underlying our technical approach.

2 Overview of Results

- *TPFDD Scheduling* – The U.S. Transportation Command and the component service commands use a relational database scheme called a TPFDD (Time-Phased Force and Deployment Data) for specifying the transportation requirements of an operation, such as Desert Storm or the Somalia relief effort. We developed a domain theory of TPFDD scheduling defining the concepts of this problem and developed laws for reasoning about them. KIDS (Kestrel Interactive Development System) was used to derive and optimize a variety of global search scheduling algorithms that perform constraint propagation [37]. The resulting code, generically called KTS (Kestrel Transportation Scheduler), has been run on a variety of TPFDDs generated by planners at USTRANSCOM and other sites.

We compared the performance of KTS with several other TPFDD scheduling systems, such as JFAST, FLOGEN, DITOPS, and PFE. We do not have access to JFAST and FLOGEN, but these are (or were) operational tools at AMC (Airlift Mobility Command, Scott AFB). According to [11] and David Brown (retired military planner consulting with the Planning Initiative), on a typical TPFDD of about 10,000 movement records, JFAST takes several hours and FLOGEN about 36 hours. KTS on a TPFDD of this size will produce a detailed schedule in *one to three minutes*. So KTS seems to be a factor of about 25 times faster than JFAST and over 250 times faster than FLOGEN. The currently operational ADANS system reportedly runs at about

the same speed as FLOGEN. KTS is orders of magnitude faster than any other TPFDD scheduler known to us.

- *Theater Scheduling*

In 1994 we began to develop a scheduler to support PACAF (Pacific Air Force) at Hickham AFB, Honolulu which is tasked with in-theater scheduling of a fleet of 26 C-130 cargo aircraft in the Pacific region. We developed (and are continuing to evolve) a theory of theater transportation scheduling. Several variants of a theater scheduler (called ITAS for In-Theater Airlift Scheduler) have been developed to date, and more are planned. The interface to ITAS and integration with a commercial database package have been developed by BBN. ITAS runs on an Apple Powerbook laptop computer. The laptop platform makes it attractive both for field and command center operations. ITAS can currently produce ATOs (Air Tasking Orders) based on the schedules that it generates.

The ITAS schedulers have emphasized flexibility and rich constraint modeling. Versions of ITAS were installed at PACAF in August 1994, September 1994, and February 1995. The most recent version simultaneously schedules the following classes of resources: (1) aircraft, (2) aircrews and their duty day cycles, (3) ground crews for unloading, and (4) ramp space at ports.

One of the reasons for the interest of PACAF in this project, is to capture some of the knowledge and experience of skilled personnel before they retire or are rotated.

- *Synthesis of constraint propagation code*

A key technical achievement of this project was discovering and implementing technology for generating efficient constraint propagation code. The speed of the KTS schedulers derives from the extremely fast checking and propagation of constraint information at every node of the runtime search tree. Whereas some knowledge-based approaches to scheduling will search a tree at the rate of several nodes per second, our synthesized schedulers search *several hundred thousand* nodes per second.

Briefly, the idea is to derive necessary conditions on feasibility of a candidate schedule. These conditions are called *cutting constraints*. The derived cutting constraints for a particular scheduling problem are analyzed to produce code that iteratively fixes violated constraints until the cutting constraints are satisfied. This iterative process subsumes the well-known processes of constraint propagation in the AI literature and the notion of cutting planes from the Operations Research literature [40, 46]. Constraint propagation is discussed in more detail in Section 5.2.3.

- *Classification approach to design*

We developed a new approach to the problem of how to construct refinements of specifications formally and incrementally. The idea is to use a taxonomy of abstract design concepts, each represented by a *design theory*. An abstract design concept is applied by constructing a specification morphism from its design theory to a requirement specification. Procedures for computing colimits and for constructing specification morphisms provide computational support for this approach. Although the classification approach applies to the incremental application of *any* kind of knowledge formally represented

in a hierarchy of theories, our work mainly focused on a hierarchy of algorithm design theories and its applications to logistical applications [39, 38]. This technique enable us to integrate at a deep semantic level problem-solving methods from Computer Science (e.g. divide-and-conquer, global search), Artificial Intelligence (e.g. heuristic search, constraint propagation, neural nets), and Operations Research (e.g. Simplex, integer programming, network algorithms). Classification is discussed in more detail in Section 6.

3 KIDS model of program development

KIDS is a program transformation system – one applies a sequence of consistency-preserving transformations to an initial specification and achieves a correct and hopefully efficient program [42]. The system emphasizes the application of complex high-level transformations that perform significant and meaningful actions. From the user’s point of view the system allows the user to make high-level design decisions like, “design a divide-and-conquer algorithm for that specification” or “simplify that expression in context”. We hope that decisions at this level will be both intuitive to the user and be high-level enough that useful programs can be derived within a reasonable number of steps.

The user typically goes through the following steps in using KIDS for program development.

1. *Develop a domain theory* – An application domain is modeled by a domain theory (a collection of types, operations, laws, and inference rules). The domain theory specifies the concepts, operations, and relationships that characterize the application and supports reasoning about the domain via a deductive inference system. Our experience has been that distributive and monotonicity laws provide most of the laws that are needed to support design and optimization of code. KIDS has a theory development component that supports the automated derivation of various kinds of laws.
2. *Create a specification* – The user enters a problem specification stated in terms of the underlying domain theory.
3. *Apply a design tactic* – The user selects an algorithm design tactic from a menu and applies it to a specification. Currently KIDS has tactics for simple problem reduction (reducing a specification to a library routine) [33], divide-and-conquer [33], global search (binary search, backtrack, branch-and-bound) [34], problem reduction generators (dynamic programming, general branch-and-bound, and game-tree search algorithms) [36], and local search (hillclimbing algorithms) [23].
4. *Apply optimizations* – The KIDS system allows the application of optimization techniques such as expression simplification, partial evaluation, finite differencing, case analysis, and other transformations [42]. The user selects an optimization method from a menu and applies it by pointing at a program expression. Each of the optimization methods are fully automatic and, with the exception of simplification (which is arbitrarily hard), take only a few seconds.

5. *Apply data type refinements* – The user can select implementations for the high-level data types in the program. Data type refinement rules carry out the details of constructing the implementation [5].
6. *Compile* – The resulting code is compiled to executable form. In a sense, KIDS can be regarded as a front-end to a conventional compiler.

Actually, the user is free to apply any subset of the KIDS operations in any order – the above sequence is typical of our experiments in algorithm design. A new system, called Specware, is currently under construction at Kestrel as a successor to KIDS. Specware is based on concepts of higher-order algebraic specifications, morphisms, and categorical constructions [20, 38, 39, 50].

4 Specifying a Scheduler

4.1 What is Scheduling?

The essential notion of scheduling is that certain activities are assigned to resources over certain time intervals. Various constraints on the assignments must be satisfied and certain measures of the cost or “goodness” of the assignment are to be optimized.

A domain theory for scheduling defines the basic concepts of scheduling and the laws for reasoning about the concepts. After a review of the relevant literature (e.g. [12]) we have identified the following general components of a scheduling domain theory.

1. *Activities* – A model of the activities can include their internal structure and characteristics, hierarchies of activity abstractions, and various operations on activities.
2. *Resources* – A model of the resources can include their internal structure and characteristics, hierarchies of resource abstractions, and various operations on resources.
3. *Time* – A time model can include a calculus of time-points or time-intervals [1, 21].
4. *Constraints* – A constraint model includes the language for stating constraints and a calculus for reasoning about them. Several classes of constraints commonly arise in practice. The most common are *precedence constraints* (which state that one activity must precede another) and *capacity constraints* (which state bounds on the capacities of resources). A constraint calculus is used to analyze constraints and to propagate the effects of new constraints through a given constraint set. Fox et al. also identify physical constraints, organizational constraints, preferences, enablement constraints, and availability constraints.
5. *Objectives* – Typically we seek to minimize the cost of a schedule. Cost can be measured in terms of time to completion, work-in-progress, total cost of consumed resources, and so on.

6. *Scheduling problem* – Using the above concepts we can formulate a variety of scheduling problems. A reservation is a triple consisting of an activity, a resource, and a time interval. Generally, a schedule is a set of *reservations* that satisfy a collection of constraints and optimize (or produce a reasonably good value of) the objective.

$$\{ \langle activity, resource, time interval \rangle \mid constraints \}.$$

Many scheduling problem instances are overconstrained – there are too few resources to schedule the activities and satisfy all constraints. Usually overconstrained problems are dealt with by relaxing the constraints and trying to satisfy as many of the constraints as possible. The usual method is to move constraints into the objective function. This entails reformulating the constraint so that it yields a quantitative measure of how well it has been satisfied. See further discussion in Section 5.2.4.

4.2 Strategic Transportation Scheduling

Transportation scheduling specializes the above general notion of scheduling: activities correspond to *movement requirements* and resources correspond to transportation assets such as planes, ships, and trucks.

A typical movement requirement has the following information

<i>move-type : movement-type</i>	↦	<i>BULK-MOVEMENT</i>
<i>quantity : integer</i>	↦	2 (STONS - Short TONS)
<i>release-date : time</i>	↦	0 (seconds from C-date)
<i>due-date : time</i>	↦	86400 (seconds from C-date)
<i>poe : port</i>	↦	<i>UHHZ</i>
<i>pod : port</i>	↦	<i>VRJT</i>
<i>distance : integer</i>	↦	5340 (nautical miles)
<i>mode : symbol</i>	↦	<i>AIR</i>

Here quantity for AIR movements is in short tons (STONS); the release and due dates are in seconds starting from C-DATE; poe (port of embarkation) and pod (port of debarkation) are given by code names; distance is in nautical miles, and the transportation mode is either AIR or SEA. A collection of movement requirements is called a TPFDD (Time-Phased Force Deployment Data).

Resources are characterized by their capacities (both passenger (PAX) and cargo capacities), and travel rate in knots.

As an example, we used a small dataset extracted from a TUNISIA TPFDD created at AFSC. This problem instance involves 480 movement requirements from 20 airports and 3 seaports to 8 airports and 2 seaports. Available air resources include KC10s, C-141s, C-5s

and sea resources include tankers (small, medium, and large), RO-ROs, LASHs, sea barges, containerhips, and breakbulks.

Eleven constraints characterize a feasible schedule for a simple TPFDD problem:

1. *Consistent POE and POD* – The POE and POD of each movement requirement on a given trip of a resource must be the same.
2. *Consistent Resource Class* – Each resource can handle only some movement types. For example, a C-141 can handle bulk and oversize movements, but not outsize movements.
3. *Consistent PAX and Cargo Capacity* – The capacity of each resource cannot be exceeded.
4. *Consistent Release Time* – The start time of a movement (its Available to Load Date (ALD)) must not precede its release time.
5. *Consistent Arrival time* – The finish time of a trip must not precede the Earliest Arrival Date (EAD) of any of the transported movement requirements.
6. *Consistent Due time* – The finish time of a movement (its Latest Arrival Date (LAD)) must not be later than its due time.
7. *Consistent Trip Separation* – Movements scheduled on the same resource must start either simultaneously or with enough separation to allow for return trips. The inherently disjunctive and relative nature of this constraint makes it more difficult to satisfy than the others.
8. *Consistent Resource Use* – Only the given resources are used.
9. *Completeness* – All movement requirements must be scheduled.

In the next section we discuss the formalization of the above concepts. This problem does not consider certain aspects of transportation scheduling, such as aircrew scheduling, ground crew scheduling, maintenance, resource utilization rates, load/unload rates, port characteristics, etc. Each of these problem features have been handled in various more elaborate specifications.

4.3 (Re-)Formulating Domain Theories for Transportation Scheduling

In the most general view, scheduling is the construction of a set of reservations that satisfy given feasibility constraints and achieve “good” values of an objective function. Formally, the schedule is a relation, or even a simple relational database. A formal domain theory based on this view is given in Appendix A. The theory provides precise definitions for the concepts, constraints, objectives, and laws used to model this application domain.

This relational view however is not always the most efficient for particular problems. We may be able to reformulate the problem, incorporating constraints and objectives, yielding a problem statement that is more amenable to efficient problem-solving. In the following we present a series of transformations that reformulate the domain theory.

In most transportation problems, each movement requirement corresponds to a unique reservation – it is scheduled exactly once with a unique resource and start time. We can make this functional dependence explicit by treating a schedule as a *map* from movement requirements to resource/time tuples. In Figure 2 we show the effect of this reformulation on the schedule datatype.

Next the trip separation constraint suggests that this map is many-to-one, since several movements can take place simultaneously on the same resource. Inverting the map will induce a partition on movement requirements. In terms of the transportation domain, inverting the map will make simultaneous movements explicit and thereby introducing the concept of a *trip* and the *manifest* of a trip.

Next we notice that the domain of a schedule map is a product of two types and these types have quite different properties (algebras): resources are a discrete set and time is (effectively) continuous and linear. The linear nature of time can be exploited by currying (to separate the two domain datatypes) and transforming the submap (from time to manifest) to a sequence, thereby making the linear structure of time explicit and introducing the concept of an *itinerary*.

This series of reformulations has dramatic effect on the trip separation constraint. In the initial formulation (in terms of reservations) this constraint involves $O(n^2)$ binary constraints between the n movements scheduled on a given resource. In the final formulation (in terms of a linearized inverse map) this constraint is reduced to $O(n)$ binary constraints between the start times of consecutive trips.

For example, on a transportation problem involving over 15,000 movement requirements obtained from the U.S. Transportation Command, the scheduler produces a complete feasible schedule in about five minutes. A straightforward constraint network formulation based on this problem data would have over 31,000 variables and 120,125,000 constraints. Incorporating some of the structure of the problem, such as the linearity of time, allows reformulating this to a system of about 108,700 constraints. However, this is still a such large formulation that it seems an implicit representation is necessary to find feasible schedules efficiently.

The final reformulation is given in Appendix B and is the theory actually used to derive a scheduler.

4.4 Formal Specification of a Scheduler

The informal specification above can be expressed as follows:

function *TS*

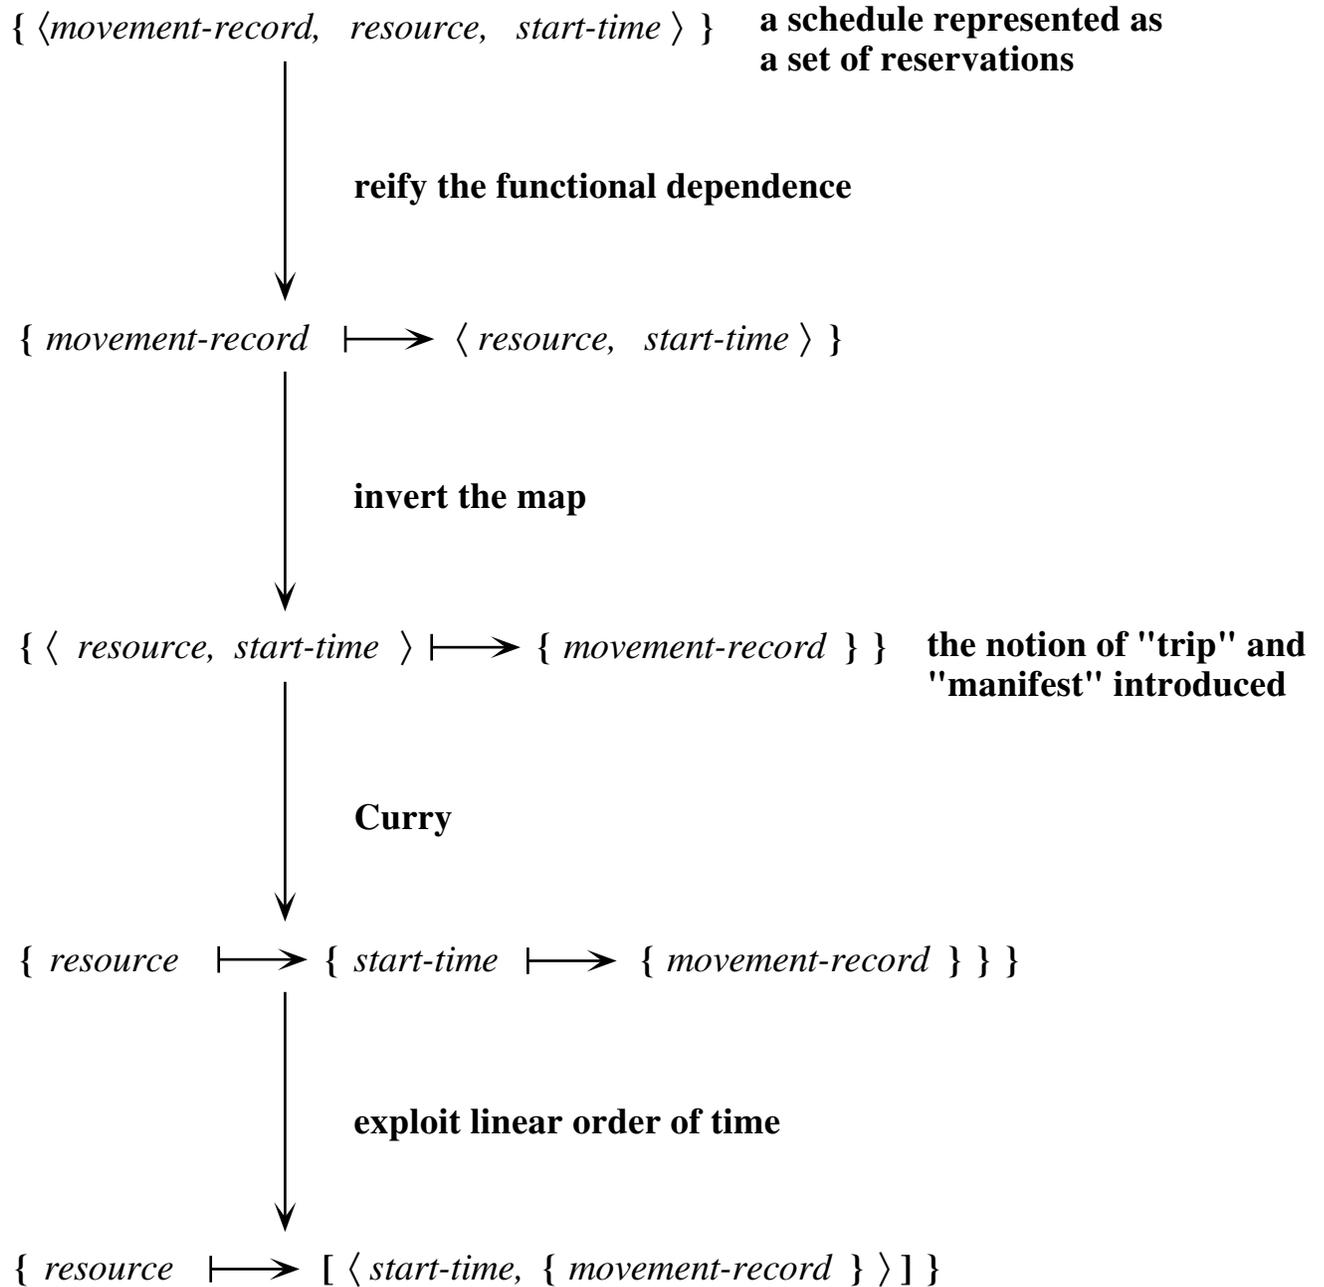


Figure 2: Reformulating a Scheduling Specification

$$\begin{aligned}
& (mvr : seq(\text{movement-record}), \\
& \quad assets : seq(\text{resource-name})) \\
\text{returns } & (sched : map(\text{resource-name}, seq(\text{trip})) \mid \\
& \quad \text{Consistent-POE}(sched) \\
& \quad \wedge \text{Consistent-POD}(sched) \\
& \quad \wedge \text{Consistent-Release-Times}(sched) \\
& \quad \wedge \text{Consistent-Arrival-Times}(sched) \\
& \quad \wedge \text{Consistent-Due-Times}(sched) \\
& \quad \wedge \text{Consistent-Trip-Separation}(sched) \\
& \quad \wedge \text{Consistent-Pax-Resource-Capacity}(sched) \\
& \quad \wedge \text{Consistent-Cargo-Resource-Capacity}(sched) \\
& \quad \wedge \text{Consistent-Movement-Type-and-Resource}(sched) \\
& \quad \wedge \text{Available-Resources-Used}(assets, sched) \\
& \quad \wedge \text{Scheduled-mvrs}(sched) = seq\text{-to-set}(mvr))
\end{aligned}$$

This specifies a function called TS that takes two inputs, a sequence of movement records called mvr s and a sequence of resources called $assets$. The function returns a schedule, which has type $map(\text{resource-name}, seq(\text{trip}))$ and must satisfy the 11 conjoined constraints. Each constraint is defined in the domain theory; for example:

function *CONSISTENT-DUE-TIMES*
 $(sched : \text{schedule}) : \text{boolean}$
 $= \forall (rsrc : \text{resource-name}, trp : \text{integer}, mvr : \text{movement-record})$
 $\quad (rsrc \in \text{domain}(sched))$
 $\quad \wedge trp \in [1..size(sched(rsrc))]$
 $\quad \wedge mvr \in sched(rsrc)(trp).manifest$
 $\quad \implies$
 $\quad sched(rsrc)(trp).start\text{-time}$
 $\quad \leq (mvr.due\text{-date} - sched(rsrc)(trp).trip\text{-duration})$

This predicate expresses the constraint that every scheduled movement-record arrives before its due date.

5 Synthesizing a Scheduler

5.1 Approach

5.1.1 Problem Theories

We briefly review some basic concepts from algebra and logic. A *theory* is a structure $\langle S, \Sigma, A \rangle$ consisting of a set of sort symbols S , operations over those sorts Σ , and axioms A

to constrain the meaning of the operations. A *theory morphism* (*theory interpretation*) maps from the sorts and operations of one theory to the sorts and expressions over the operations of another theory such that the image of each source theory axiom is valid in the target theory. A *parameterized theory* has formal parameters that are themselves theories [15]. The binding of actual values to formal parameters is accomplished by a theory morphism. Theory $\mathcal{T}_2 = \langle S_2, \Sigma_2, A_2 \rangle$ *extends* (or is an *extension* of) theory $\mathcal{T}_1 = \langle S_1, \Sigma_1, A_1 \rangle$ if $S_1 \subseteq S_2$, $\Sigma_1 \subseteq \Sigma_2$, and $A_1 \subseteq A_2$.

Problem theories define a problem by specifying a domain of problem instances or inputs and the notion of what constitutes a solution to a given problem instance. Formally, a *problem theory* \mathcal{B} has the following structure.

Sorts	D, R
Operations	$I : D \rightarrow \text{Boolean}$
	$O : D \times R \rightarrow \text{Boolean}$

The *input condition* $I(x)$ constrains the input domain D . The *output condition* $O(x, z)$ describes the conditions under which output domain value $z \in R$ is a *feasible solution* with respect to input $x \in D$. Theories of booleans and sets are implicitly imported. Problems of finding optimal feasible solutions can be treated as extensions of problem theory by adding a cost domain, cost function, and ordering on the cost domain.

For example, the problem of finding feasible schedules can be presented as a problem theory via a theory interpretation into the domain theory of transportation scheduling:¹

D	\mapsto	$seq(\text{movement-record}) \times seq(\text{resource})$
I	\mapsto	$\lambda(Mvrs, resources) \text{ true}$
R	\mapsto	$map(\text{resource}, seq(\text{trip}))$
O	\mapsto	$\lambda(Mvrs, resources, sched)$
		$Consistent-POE(sched)$
		$\wedge Consistent-POD(sched)$
		$\wedge Consistent-Release-Times(sched)$
		$\wedge Consistent-Arrival-Times(sched)$
		$\wedge Consistent-Due-Times(sched)$
		$\wedge Consistent-Trip-Separation(sched)$
		$\wedge Consistent-Pax-Resource-Capacity(sched)$
		$\wedge Consistent-Cargo-Resource-Capacity(sched)$
		$\wedge Consistent-Movement-Type-and-Resource(sched)$
		$\wedge Available-Resources-Used(resources, sched)$
		$\wedge Scheduled-mvrs(sched) = seq-to-set(mvrs)$

¹The domain theory includes definitions for the types of movement-record, resource, trip (a record comprised of start-time and manifest), and schedule (a map from resource to sequence of trip).

5.1.2 Algorithm Theories

An *algorithm theory* represents the essential structure of a certain class of algorithms \mathcal{A} [43]. Algorithm theory \mathcal{A} extends problem theory \mathcal{B} with any additional sorts, operators, and axioms needed to support the correct construction of an \mathcal{A} algorithm for \mathcal{B} . A theory morphism from the algorithm theory into some problem domain theory provides the problem-specific concepts needed to construct an instance of an \mathcal{A} algorithm.

For example, global search theory (presented below in Section 6.2.1) extends problem theory with the basic concepts of backtracking: subspace descriptors, initial space, the splitting and extraction operations, filters, and so on. A divide-and-conquer theory would extend problem theory with concepts such as decomposition operators and composition operators [33, 36].

5.2 Synthesizing a Scheduler

There are two basic approaches to computing a schedule: local and global. Local methods focus on individual schedules and similarity relationships between them. Once an initial schedule is obtained, it is iteratively improved by moving to neighboring structurally similar schedules. Repair strategies [53, 25, 4, 31], and fixpoint iteration [8], and linear programming algorithms are examples of local methods.

Global methods focus on sets of schedules. A feasible or optimal schedule is found by repeatedly splitting an initial set of schedules into subsets until a feasible or optimal schedule can be easily extracted. Backtrack, heuristic search, and branch-and-bound methods are all examples of global methods. We explore the application of global methods. In the following subsections we formalize the notion of global search method and show how it can be applied to synthesize a scheduler. Other projects taking a global approach include ISIS [13], OPIS/DITOPS [47], and MicroBoss [28] (all at CMU).

5.2.1 Global Search Theory

The basic idea of global search is to represent and manipulate sets of candidate solutions. The principal operations are to *extract* candidate solutions from a set and to *split* a set into subsets. Derived operations include various *filters* which are used to eliminate sets containing no feasible or optimal solutions. Global search algorithms work as follows: starting from an initial set that contains all solutions to the given problem instance, the algorithm repeatedly extracts solutions, splits sets, and eliminates sets via filters until no sets remain to be split. The process is often described as a tree (or DAG) search in which a node represents a set of candidates and an arc represents the split relationship between set and subset. The filters serve to prune off branches of the tree that cannot lead to solutions.

The sets of candidate solutions are often infinite and even when finite they are rarely represented extensionally. Thus global search algorithms are based on an abstract data type

of intensional representations called *space descriptors* (denoted by hatted symbols). In addition to the extraction and splitting operations mentioned above, the type also includes a predicate *satisfies* that determines when a candidate solution is in the set denoted by a descriptor. Further, there is a refinement relation on spaces that corresponds to the subset relation on the sets denoted by a pair of descriptors.

The various operations in the abstract data type of space descriptors together with problem specification can be packaged together as a theory. Formally, abstract *global search theory* (or simply *gs-theory*) \mathcal{G} is presented in Figure 3, where D is the input domain, R is the output domain, I is the input condition, O is the output condition, \hat{R} is the type of space descriptors, \hat{I} defines legal space descriptors, \hat{r} and \hat{s} vary over descriptors, $top(x)$ is the descriptor of the initial set of candidate solutions, $Satisfies(z, \hat{r})$ means that z is in the set denoted by descriptor \hat{r} or that z satisfies the constraints that \hat{r} represents, and $Extract(z, \hat{r})$ means that z is directly extractable from \hat{r} .

The relations *Split-Arg* and *Split-Constraint* are used to determine and perform splitting. In particular, if $Split-Arg(x, \hat{r}, c)$ then c is information that characterizes (or informs) one branch of the split. $Split-Constraint(x, \hat{r}, c, \hat{s})$ means that \hat{s} results from incorporating information c into the descriptor \hat{r} (with respect to input x). *Split-Arg* is used to control the generation of children of a node in the search tree and *Split-Constraint* is used to specify one child. *Split-Constraint* can be thought of as a parameterized constraint whose alternative arguments are supplied by *Split-Arg*.

The refinement relation $\hat{r} \sqsupseteq \hat{s}$ holds when \hat{s} denotes a subset of the set denoted by \hat{r} . Further, \hat{R} together with \sqsupseteq forms a bounded semilattice. This structure will play a crucial role in constraint propagation algorithms.

Note that all variables in the axioms are assumed to be universally quantified unless explicitly specified otherwise. Axiom GS0 asserts that the initial descriptor $\hat{r}_0(x)$ is a legal descriptor. Axiom GS1 asserts that legal descriptors split into legal descriptors and that *Split* induces a well-founded ordering on spaces. Axiom GS2 constrains the denotation of the initial descriptor — all feasible solutions are contained in the initial space. Axiom GS3 gives the denotation of an arbitrary descriptor \hat{r} — an output object z is in the set denoted by \hat{r} if and only if z can be extracted after finitely many applications of *Split* to \hat{r} where

$$Split^*(x, \hat{r}, \hat{s}) \iff \exists(k : Nat) \ Split^k(x, \hat{r}, \hat{s})$$

and

$$Split^0(x, \hat{r}, \hat{t}) \iff \hat{r} = \hat{t}$$

and for all natural numbers k

$$Split^{k+1}(x, \hat{r}, \hat{t}) \iff \exists(\hat{s} : \hat{R}, i : \hat{C}) (Split-Arg(x, \hat{r}, i) \wedge Split-Constraint(x, \hat{r}, i, \hat{s}) \wedge Split^k(x, \hat{s}, \hat{t})).$$

Axiom GS4 asserts that if \hat{r} splits to \hat{s} then \hat{r} also refines to \hat{s} ; thus the refinement relation on \hat{R} is weaker than the split relation. We also need the axioms that $\langle \hat{R}, \sqsupseteq, \sqcap \rangle$ is a semilattice. For simplicity, we write $\hat{r} \sqsupseteq \hat{s}$ rather than the correct $\sqsupseteq(x, \hat{r}, \hat{s})$; and similarly $\hat{r} \sqcap \hat{s}$.

Spec *Global-Search*

Sorts	D	<i>input domain</i>
	R	<i>output domain</i>
	\hat{R}	<i>subspace descriptors</i>
	\hat{C}	<i>splitting information</i>

Operations

$I : D \rightarrow \text{boolean}$	<i>input condition</i>
$O : D \times R \rightarrow \text{boolean}$	<i>input/output condition</i>
$\hat{I} : D \times \hat{R} \rightarrow \text{boolean}$	<i>subspace descriptors condition</i>
$Satisfies : R \times \hat{R} \rightarrow \text{boolean}$	<i>denotation of descriptors</i>
$Split\text{-}Arg : D \times \hat{C} \times \hat{R} \rightarrow \text{boolean}$	<i>specifies arguments to split constraint</i>
$Split\text{-}Constraint : D \times \hat{R} \times \hat{C} \times \hat{R} \rightarrow \text{boolean}$	<i>parameterized splitting constraint</i>
$Extract : R \times \hat{R} \rightarrow \text{boolean}$	<i>extractor of solutions from spaces</i>
$\Psi : D \times R \times \hat{R} \rightarrow \text{boolean}$	<i>cutting constraint</i>
$\xi : D \times \hat{R} \rightarrow \text{boolean}$	<i>cutting constraint</i>
$\sqsupseteq : D \times \hat{R} \times \hat{R} \rightarrow \text{boolean}$	<i>refinement relation</i>
$top : D \rightarrow \hat{R}$	<i>initial space</i>
$bot : \hat{R}$	<i>inconsistent space</i>

Axioms

- GS0. All feasible solutions are in the *top* space
$$I(x) \wedge O(x, z) \implies Satisfies(z, top(x))$$
- GS1. All solutions in a space are finitely extractable
$$I(x) \wedge \hat{I}(x, \hat{r}) \implies (Satisfies(z, \hat{r}) \iff \exists(\hat{s}) (Split^*(x, \hat{r}, \hat{s}) \wedge Extract(z, \hat{s})))$$
- GS2. Specification of Cutting Constraint
$$Satisfies(z, \hat{r}) \wedge O(x, z) \implies \Psi(x, z, \hat{r})$$
- GS3. Definition of Cutting Constraint on Spaces
$$\xi(x, \hat{r}) \iff \forall(z : R) (Sat(z, \hat{r}) \implies \Psi(x, z, \hat{r}))$$
- GS4. Definition of Refinement
$$\hat{r} \sqsupseteq \hat{s} \iff \forall(z : R) (Satisfies(z, \hat{s}) \implies Satisfies(z, \hat{r}))$$
- GS5. $\langle \hat{R}, \sqsupseteq, \sqcap, top, bot \rangle$ is a bounded meet-semilattice with *bot* as universal lower bound.

end spec

Figure 3: Global Search Theory

For example, a simple global search theory of scheduling has the following form. Schedules are represented as maps from resources to sequences of trips, where each trip includes earliest-start-time, latest-start-time, port of embarkation, port of debarkation, and manifest (set of movement records or ULNs + CINs + PINs from the TPFDD). The type of schedules has the invariant (or subtype characteristic) that for each trip, the earliest-start-time is no later than the latest-start-time. A partial schedule is a schedule over a subset of the given movement records.

The initial (partial) schedule is just the empty schedule – a map from the available resources to the empty sequence of trips. A partial schedule is extended by first selecting a movement record mvr to schedule, then selecting a resource r , and then a trip t on r (either an existing trip or a newly created one) – the triple $\langle mvr, r, t \rangle$ constitutes the information c of *Split-Arg. Split-Constraint* given $\langle mvr, r, t \rangle$ creates an extended schedule that has mvr added to the manifest of trip t on resource r . The alternative ways that a partial schedule can be extended naturally gives rise to the branching structure underlying global search algorithms.

The formal version of this global search theory of scheduling can be inspected in the domain theory in Appendix C.

5.2.2 Pruning Mechanisms

When a partial schedule is extended it is possible that some problem constraints are violated in such a way that further extension to a complete feasible schedule is impossible. In tree search algorithms it is crucial to detect such violations as early as possible.

Pruning tests are derived in the following way. The test

$$\exists(z) (Satisfies(z, \hat{r}) \wedge O(x, z)) \tag{1}$$

decides whether there exist any feasible solutions that are in the space denoted by \hat{r} . If we could decide this at each node of our branching structure then we would have perfect search – no deadend branches would ever be explored. In practice it would be impossible or horribly complex to compute (1), so we rely instead on an inexpensive approximation to it. In fact, if we approximate (1) by weakening it (deriving a necessary condition of it) we obtain a sound pruning test. That is, suppose we can derive a test $\Phi(x, \hat{r})$ such that

$$\exists(sched) (Satisfies(z, \hat{r}) \wedge O(x, z)) \implies \Phi(x, \hat{r}). \tag{2}$$

By the contrapositive of (2), if $\neg\Phi(x, \hat{r})$ then there are no feasible solutions in \hat{r} , so we can eliminate it from further processing. A global search algorithm will test Φ at each node it explores, pruning those nodes where the test fails.

More generally, necessary conditions on the existence of feasible (or optimal) solutions below a node in a branching structure underlie pruning in backtracking and the bounding and dominance tests of branch-and-bound algorithms [34].

It appears that the bottleneck analysis advocated in the constraint-directed search projects at CMU [12, 28] leads to a semantic approximation to (1), but neither a necessary nor sufficient

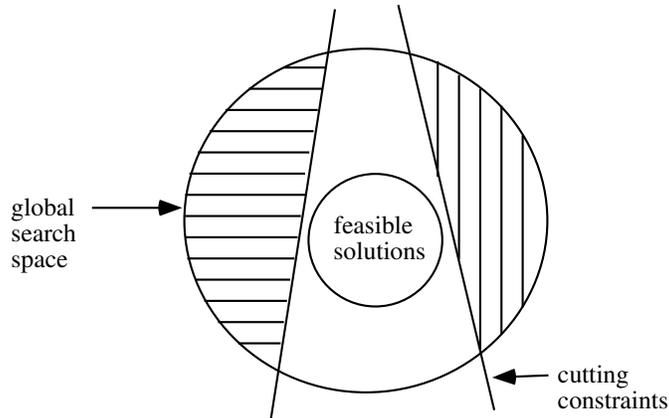


Figure 4: Global Search Subspace and Cutting Constraints

condition. Such a *heuristic* evaluation of a node is inherently fallible, but if the approximation is close enough it can provide good search control with relatively little backtracking.

To derive pruning tests for the strategic transportation scheduling problem, we instantiate (1) with our definition of *Satisfies* and O and use an inference system to derive necessary conditions. The resulting tests are fairly straightforward; of the 11 original feasibility constraints, 7 yield pruning tests on partial schedules. For example, the partial schedule must satisfy *Consistent-POE*, *Consistent-POD*, *Consistent-Pax-Resource-Capacity*, *Consistent-Cargo-Resource-Capacity*, *Consistent-Movement-Type-and-Resource*, and *Available-Resources-Used*. The reader may note that computing these tests on partial schedules is rather expensive and mostly unnecessary – later program optimization steps will however reduce these tests to fast and irredundant form. For example, the first test will reduce to checking that when we place a movement record *mvr* on trip t , we check that the POE of *mvr* and t are consistent.

For details of deriving pruning mechanisms for other problems see [34, 42, 43, 35].

5.2.3 Cutting Constraints and Constraint Propagation

Constraint propagation is a more general technique that is crucial for early detection of infeasibility. We developed a general mechanism for deriving constraint propagation code and applied it to scheduling.

Each node in a backtrack tree can be viewed as a data structure that denotes a set of candidate solutions – in particular the solutions that occur in the subtree rooted at the node (see Figure 4). Thus the root denotes the set of all candidate solutions found in the tree.

Pruning has the effect of removing a node (set of solutions) from further consideration. In contrast, constraint propagation has the effect of changing the space descriptor so that it

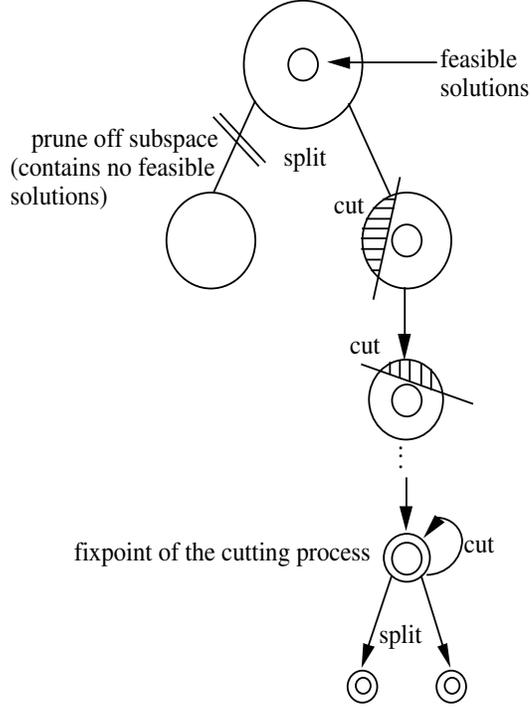


Figure 5: Pruning and Constraint Propagation

denotes a smaller set of candidate solutions. The effect of constraint propagation is to spread information through the subspace descriptor resulting in a tighter descriptor and possibly exposing infeasibility. Pruning can be treated as a special case of propagation in which a space is refined to descriptor that denotes the empty set of solutions.

Constraint propagation is based on the notion of *cutting constraints* which are necessary conditions $\Psi(x, z, \hat{r})$ that a candidate solution z satisfying \hat{r} is feasible:

$$\forall(x : D, \hat{r} : \hat{R}, z : R)(Satisfies(z, \hat{r}) \wedge O(x, z) \implies \Psi(x, z, \hat{r})) \quad (3)$$

See Figures 4 and 5. In order to get a test on spaces that decides whether Ψ has been incorporated, we make one further definition:

$$\xi(x, \hat{r}) \iff \forall(z : R)(Satisfies(z, \hat{r}) \implies \Psi(x, z, \hat{r})) \quad (4)$$

The test $\xi(x, \hat{r})$ holds exactly when all candidate solutions in \hat{r} satisfy Ψ , and we say that \hat{r} *satisfies* ξ .

The key question at this point is: Given a descriptor \hat{r} that doesn't satisfy ξ , how can we incorporate ξ into \hat{r} ? The answer is to find the greatest refinement of \hat{r} that satisfies ξ ; we say \hat{t} *incorporates* ξ into \hat{r} if

$$\hat{t} = \max_{\sqsupseteq} \{\hat{s} \mid \hat{r} \sqsupseteq \hat{s} \wedge \xi(x, \hat{s})\}. \quad (5)$$

which asserts that \hat{t} is maximal over the set of descriptors that refine \hat{s} and satisfy ξ , with respect to ordering \sqsupseteq . We want \hat{t} to be a refinement of \hat{r} so that all of the information in \hat{r} is preserved and we want \hat{t} to be maximal so that no other information than \hat{r} and ξ is incorporated into \hat{t} .

The next question concerns the conditions under which Formula (5) is satisfiable. Assuming that \hat{R} is a semilattice, we can use variants of Tarski's fixpoint theorem (c.f. [8]):

Theorem If there is a function f such that

1. f is monotonic on \hat{R} (i.e. $\hat{s} \sqsupseteq \hat{t} \implies f(x, \hat{s}) \sqsupseteq f(x, \hat{t})$)
2. f is deflationary (i.e. $\hat{r} \sqsupseteq f(x, \hat{r})$)
3. f has fixed-points satisfying ξ (i.e. $f(x, \hat{r}) = \hat{r} \iff \xi(x, \hat{r})$)

then (1) $\hat{t} = \max_{\sqsupseteq} \{\hat{s} \mid \hat{r} \sqsupseteq \hat{s} \wedge \xi(x, \hat{s})\}$ exists

and (2) \hat{t} is the greatest fixpoint of f ; i.e. \hat{t} can be computed by iteratively applying f to \hat{r} until a fixpoint is reached.

The challenge is to construct a monotonic, deflationary function whose fixed-points satisfy ξ . A general construction in terms of global search theory can be sketched as follows. Let

$$f(x, \hat{r}) = \begin{cases} \hat{r} & \text{if } \xi(x, \hat{r}) \\ \dots & \text{if } \neg\xi(x, \hat{r}) \end{cases}$$

The intent is to define f so that it has fixpoints exactly when $\xi(x, \hat{r})$ holds. When $\xi(x, \hat{r})$ doesn't hold, then we know (by the definition of ξ and the contrapositive of formula (3)) that

$$\exists(z : R)(\text{Satisfies}(z, \hat{r}) \wedge \neg O(x, z))$$

i.e. there are some infeasible solutions in the space described by \hat{r} . Ideally $\neg\xi(x, \hat{r})$ is a *constructive* assertion, so it provides information on *which* solutions are infeasible and how to eliminate them. In place of the ellipsis above we require a new descriptor that refines \hat{r} (so f is decreasing on all inputs), allows f to be monotone, and eliminates some of the infeasible solutions indicated by $\neg\xi(x, \hat{r})$. In general it is difficult to see how to achieve this end without assuming special structure to \hat{R} and ξ .

We have identified some special cases for which an analytic procedure can produce the necessary iteration function f from ξ . These special cases subsume our scheduling applications and many related Constraint Satisfaction Problems (CSP) problems. Suppose that the constraint ξ has the form

$$B(x, \hat{r}) \sqsupseteq \hat{r} \quad (6)$$

where $B(x, \hat{r})$ is monotonic in \hat{r} . We say that ξ is a *Horn-like constraint* by generalization of Horn clauses in logic. Notice that the occurrence of \hat{r} on the right-hand side of the inequality has positive polarity (i.e. it is monotonic in \hat{r}), whereas the occurrence(s) of \hat{r} on the left-hand side have negative polarity (i.e. are antimonotonic). If the constraint were boolean (with B and \hat{r} being boolean values and \sqsupseteq being implication), then this would be called a definite Horn clause. When our constraints are Horn-like, then there is a simple definition for the desired function f :

$$f(x, \hat{r}) = \begin{cases} \hat{r} & \text{if } B(x, \hat{r}) \sqsupseteq \hat{r} \\ B(x, \hat{r}) \sqcap \hat{r} & \text{if } \neg B(x, \hat{r}) \sqsupseteq \hat{r} \end{cases}$$

or equivalently

$$f(x, \hat{r}) = B(x, \hat{r}) \sqcap \hat{r}.$$

It is easy to check that f is monotone in \hat{r} , deflationary, and has fixed-points exactly when ξ holds. Therefore, simple iteration of f will converge to the descriptor that incorporates ξ into \hat{r} . However, if \hat{r} is an aggregate structure such as a tuple or map, then the changes made at each iteration may be relatively sparse, so the simple iteration approach may be grossly inefficient. We found this feature to be characteristic of scheduling and other CSPs. Our approach to solving this problem is to focus on single point changes and to exploit dependence analysis. For each component of \hat{r} we define a separate change propagation procedure. The arguments to a propagation procedure specify a change to the component. This change is performed and then the change procedures for all other components that could be affected by the change are invoked. Static dependence analysis at design-time is used to determine which constraints could be affected by a change to a given component.

A program scheme for global search with constraint propagation is presented in Figure 6. The global search design tactic in KIDS first instantiates this scheme, then invokes a tactic for synthesizing propagation code to satisfy the specification *F-split-and-propagate*.

CSPs with Horn-like constraints

We now elaborate the previous exposition of propagation of Horn-like constraints arising in CSPs. To keep matters simple, yet general, suppose that the output datatype R is $map(VAR, VALSET)$, where VAR is a type of variables, and $VALSET$ is a type that denotes a set of values (this implies that all the variables have the same type and refinement ordering), and the \sqsupseteq relation has the form:

$$\hat{r} \sqsupseteq \hat{s} \text{ iff } \bigwedge_v \hat{r}(v) \sqsupseteq \hat{s}(v).$$

Suppose further that ξ is a conjunction of constraints giving bounds on the variables:

$$\xi(x, \hat{r}) \iff \bigwedge_v B_v(x, \hat{r}) \sqsupseteq \hat{r}(v)$$

Spec *Global-Search-Program* ($T :: \text{Global-Search}$)

Operations

F-initial-propagate ($x : D \mid I(x)$)

returns ($\hat{t} : \hat{R} \mid \hat{t} = \max_{\sqsupseteq} \{\hat{s} \mid \text{top}(x) \sqsupseteq \hat{s} \wedge \hat{I}(x, \hat{s}) \wedge \xi(x, \hat{s})\}$)

F-split-and-propagate

($x : D, \hat{r} : \hat{R}, c : \hat{C}$)

$\mid I(x) \wedge \hat{I}(x, \hat{r}) \wedge \text{Split-Arg}(x, \hat{r}, c) \wedge \xi(x, \hat{r}) \wedge \hat{r} \neq \text{bot}$)

returns ($\hat{t} : \hat{R} \mid \hat{t} = \max_{\sqsupseteq} \{\hat{s} \mid \hat{r} \sqsupseteq \hat{s} \wedge \hat{I}(x, \hat{s})$
 $\wedge \text{Split}(x, \hat{r}, c, \hat{s}) \wedge \xi(x, \hat{s})\}$)

F-gs ($x : D, \hat{r} : \hat{R} \mid I(x) \wedge \hat{I}(x, \hat{r}) \wedge \Phi(x, \hat{r})$)

returns ($z : R \mid O(x, z) \wedge \text{Satisfies}(z, \hat{r})$)

= **if** $\exists(z) (\text{Extract}(z, \hat{r}) \wedge O(x, z))$

then $\text{some}(z) (\text{Extract}(z, \hat{r}) \wedge O(x, z))$

else $\text{some}(z) \exists(c : \hat{C}, \hat{t} : \hat{R})$

$(\text{Split-Arg}(x, \hat{r}, c)$

$\wedge \hat{t} = \text{F-split-and-propagate}(x, \hat{r}, c) \wedge \hat{t} \neq \text{bot}$

$\wedge z = \text{F-gs}(x, \hat{t})$)

F ($x : D \mid I(x)$)

returns ($z : R \mid O(x, z)$)

= $\text{some}(z) \exists(\hat{t}) (\hat{t} = \text{F-initial-propagate}(x)$

$\wedge \hat{t} \neq \text{bot}$

$\wedge z = \text{F-gs}(x, \hat{t})$)

end spec

Figure 6: Global Search Program Theory

where $B_v(x, \hat{r})$ is monotonic in \hat{r} . Under these assumptions, $\neg\xi(x, \hat{r})$ implies that the bounding constraint on some variable v is violated; i.e.

$$\neg B_v(x, \hat{r}) \sqsupseteq \hat{r}(v).$$

To “fix” such a violation we can change the current valset of v to

$$B_v(x, \hat{r}) \sqcap \hat{r}(v),$$

which simultaneously refines $\hat{r}(v)$, since

$$\hat{r}(v) \sqsupseteq B_v(x, \hat{r}) \sqcap \hat{r}(v)$$

and reestablishes the constraint on v , since

$$B_v(x, \hat{r}) \sqsupseteq B_v(x, \hat{r}) \sqcap \hat{r}(v).$$

Let

$$B(x, \hat{r}) = \{ | u \rightarrow B_u(x, \hat{r}) \sqcap \hat{r}(u) \mid u \in \text{domain}(\hat{r}) \}$$

then, define f as:

$$f(x, \hat{r}) = \hat{r} \sqcap B(x, \hat{r})$$

Constraint propagation is treated here as iteration of f until a fixed-point is reached. Efficiency requires that we go farther, since only a sparse subset of the variables in \hat{r} will be updated at each iteration. If we implemented the iteration on a vector processor or SIMD machine, the overall computation could be fast, but wasteful of processors. On a sequential machine, it is advantageous to analyze the constraints in ξ to infer dependence of constraints on variables. That is, if (the valset of) variable v changes, which constraints in ξ could become violated? This dependence analysis can be used to generate special-purpose propagation code as follows.

For each variable v , let $\text{affects}(v)$ be the set of variables whose constraints could be violated by a change in v ; more formally, let

$$\text{affects}(v) = \{u \mid v \text{ occurs in } B_u\}.$$

We can then generate a set of procedures that carry out the propagation/iteration of f : For each variable v , generate the following propagation procedure:

```

Propagatev (x : D,  $\hat{r}$  :  $\hat{R}$ , new-valset : VALSET
  | I(x)  $\wedge$   $\hat{I}(x, \hat{r})$ 
   $\wedge$   $\hat{r}(v) \sqsupseteq$  new-valset
   $\wedge$   $B_v(x, \hat{r}) \sqsupseteq$  new-valset)
= let ( $\hat{s}$  :  $\hat{R}$  = map-shadow( $\hat{r}$ , v, new-valset))
  if  $\neg\hat{I}(x, \hat{s})$  then bot
  else
  ... for each variable  $u$  in  $\text{affects}(v)$  ...

```

```

... generate the following code block ...
if  $\hat{s} = bot$  then  $bot$ 
else (if  $\neg(B_u(x, \hat{s}) \sqsupseteq \hat{s}(u))$ 
      then  $\hat{s} \leftarrow Propagate_u(x, \hat{s}, B_u(x, \hat{s}) \sqcap \hat{s}(u))$ );
...
 $\hat{s}$ 
end

```

where $map-shadow(\hat{r}, v, new-valset)$ returns the map \hat{r} modified so that $\hat{r}(v) = new-valset$.

To finish up, if $Split(x, \hat{r}, i, \hat{s})$ has the form

$$\hat{s}(u) = C(x, \hat{r}, i)$$

for some function C that yields a refined valset for variable u , then we can satisfy *F-split-and-propagate* as follows:

$$F-split-and-propagate(x, \hat{r}, i) = propagate_u(\hat{r}, C(x, \hat{r}, i)).$$

The change to u induced in the call to $propagate_u$ will in turn trigger changes to other variables, and so on.

We have described the generation of constraint propagation in a relatively simple setting. One of the authors (Westfold) was largely responsible for the development and implementation of this work. The implementation treats a much broader range of problem features than has been described above. Further elaborations include

1. Heterogeneous variables (and semilattice/refinement structure)
2. Multiple constraints on each variable
3. Indexed variables
4. Conditional constraints
5. Dynamic set of variables
6. Ordering of constraints in propagation procedures

There are many ways to implement constraint propagation, this being just one. Our approach is useful when the *affects* relation is relatively sparse, so special control code to follow dependences and fixing violations is efficient. An alternative approach is to reify *affects* via explicit links between variables, forming a constraint network. The synthesis of the propagation control strategy is relatively simple, since we only need to follow dependence links. Disadvantages of this approach include the size of the constraint network and the cost of maintaining it. This is a common approach in the CSP literature.

Our model of constraint propagation generalizes the concepts of cutting planes in the Operations Research literature [26] and the forms of propagation studied in the constraint satisfaction literature (e.g. [18]). Our use of fixed-point iteration for constraint propagation is similar to Paige’s work on fixed-point iteration in RAPTS [8]. The main differences are (1) RAPTS expects the user to supply the monotone function as part of the initial specification whereas we derive it from a more abstract statement of the problem; (2) RAPTS instantiates a straightforward iteration scheme and then performs optimizations. Such an approach would be too inefficient for scheduling applications, so we use dependence analysis to generate code that is specific to the constraint system at hand.

Constraint Propagation for Transportation Scheduling

For transportation scheduling, each iteration of the *Propagate* operation has the following form, where est_i denotes the earliest-start-time for trip i and est'_i denotes the next value of the earliest-start-time for trip i (analogously, lst_i denotes latest-start-time), and $roundtrip_i$ is the roundtrip time for trip i on resource r . For each resource r and the i^{th} trip on r ,

$$est'_i = \max \begin{cases} est_i \\ est_{i-1} + roundtrip_i \\ \text{max-release-time}(\text{manifest}_i) \end{cases}$$

$$lst'_i = \min \begin{cases} lst_i \\ lst_{i+1} - roundtrip_i \\ \text{min-finish-time}(\text{manifest}_i) \end{cases}$$

Here $\text{max-release-time}(\text{manifest}_i)$ computes the max over all of the release times of movement requirements in the manifest of trip i and $\text{min-finish-time}(\text{manifest}_i)$ computes the minimum of the finish times of movement requirements in the same manifest. Boundary cases must be handled appropriately.

After adding a new movement record to some trip, the effect of *Propagate* will be to shrink the

$$\langle \text{earliest-start-time}, \text{latest-start-time} \rangle$$

window of each trip on the same resource. If the window becomes negative for any trip, then the partial schedule is necessarily infeasible and it can be pruned.

The constraint propagation code generated for *TS* in Appendix D is nearly as fast as hand-written propagation code for the same problem (cf. Appendix C in [37]).

5.2.4 Constraint Relaxation

Many scheduling problems are overconstrained. Overconstrained problems are typically handled by relaxing the constraints. The usual method, known as Lagrangian Relaxation [26], is to move constraints into the objective function. This entails reformulating the constraint so that it yields a quantitative measure of how well it has been satisfied.

Another approach is to relax the input data just enough that a feasible solution exists. To test this approach, we hand-modified one version of KTS so it relaxes the LAD (Latest Arrival Date) constraint. The relaxation takes place only when there is no feasible solution to the problem data. KTS keeps track of a quantitative measure of each LAD violation (e.g. the difference between the arrival date of a trip and the LAD of a movement requirement in that trip). If there is no feasible reservation for the movement requirement being scheduled, then KTS uses the recorded information to relax its the LAD. The relaxation is such as to minimally delay the arrival of the requirement to its POD.

This technique, which we call *data relaxation*, can be described more generally. Suppose that we specify a certain constraint to be relaxable. Whenever we detect that the input data has no feasible solution, we attempt to relax the input data just enough to allow a feasible solution. Of course, the problem-solving process and data relaxation are interleaved.

At each global search iteration we evaluate this objective function for all candidate solutions. Using these values the algorithm takes a greedy decision of which branch of the global search tree should be split next. The result is a heuristically-guided algorithm that finds good but not necessarily optimal schedules.

It remains an open task to formalize the notion of data relaxation and to develop a tactic for synthesizing relaxation code in the context of global search with constraint propagation.

5.2.5 Using KIDS

In developing a new scheduling application, most of the user's time is spent building a theory of the domain. Our scheduling theories have evolved over months of effort into 50-70 pages of text. It currently takes about 90 minutes to transform our most complex scheduling specification (for ITAS) into optimized and compiled CommonLisp code for Sun workstations. Evolution of the scheduler is performed by evolving the domain theory and specification, followed by regeneration of code.

Currently, the global search design tactic in KIDS is used to design an algorithm for F and F - gs in Figures-scheme. A specialized tactic for generating constraint propagation code for Horn-like constraints is used to generate code for *F-split-and-propagate*. Once the algorithm is designed, then a series of simplification and common-subexpression-elimination transformations are applied. A trace of the KIDS derivation is given in Appendix D. See [42] for a detailed description of a session with KIDS.

Data Sets (Air only)	# of input TPFDD records (ULNs)	# of individual movements	# of scheduled items after splitting	Solution time	Msec per scheduled item
CDART		296	330	0.5 sec	1.5
CSRT01	1,600	1,261	3,557	44 sec	12
096-KS	20,400	4,644	6,183	86 sec	14
9002T Borneo	28,900	10,623	15,119	290 sec	20

Figure 7: KTS Scheduling Statistics

5.3 KTS – Strategic Transportation Scheduling

The KTS schedulers synthesized using the KIDS program transformation system are extremely fast and accurate [44, 45]. The chart in Figure 7 lists 4 TPFDD problems, and for each problem (1) the number of TPFDD lines (each requirement line contains up to several hundred fields), (2) the number of individual movement requirements obtained from the TPFDD line (each line can specify several individual movements requirements), (3) the number of movement requirements obtained after splitting (some requirements are too large to fit on a single aircraft or ship so they must be split), (4) the cpu time to generate a complete schedule, and (5) time spent per scheduled movement. Similar results were obtained for sea movements. The largest problem, Borneo NEO, is harder to solve, because of the presence of 29 movement requirements that are inherently unschedulable: their due date comes before their availability date. Such inconsistencies must be expected and handled by a realistic system. KTS simply relaxes the due date the minimal amount necessary to obtain a feasible schedule.

We compared the performance of KTS with several other TPFDD scheduling systems: JFAST, FLOGEN, DITOPS, and PFE. We do not have access to JFAST and FLOGEN, but these are (or were) operational tools at AMC (Airlift Mobility Command, Scott AFB). According to [11] and David Brown (retired military planner consulting with the Planning Initiative), on a typical TPFDD of about 10,000 movement records, JFAST takes several hours and FLOGEN about 36 hours. KTS on a TPFDD of this size will produce a detailed schedule in *one to three minutes*. So KTS seems to be a factor of about 25 times faster than JFAST and over 250 times faster than FLOGEN. The currently operational ADANS system reportedly runs at about the same speed as FLOGEN. When comparing schedulers it is also important to compare the transportation models that they support. KTS has a richer model than JFAST (i.e. handles more constraints and problem features), but less rich than ADANS. The ITAS effort described in the next section reflects our efforts to synthesize

schedulers that have at least the richness of the ADANS model.

The DITOPS project at CMU also models scheduling as a constraint satisfaction problem. However, DITOPS effectively interprets its problem constraints, whereas the transformational approach can produce highly optimized “compiled” constraint operations. DITOPS emphasizes complex heuristics for guiding the search away from potential bottlenecks. In contrast KTS uses simple depth-first search but emphasizes the use of strong and extremely fast pruning and constraint propagation code. DITOPS requires minutes to solve the CDART data. KTS finds a complete feasible solution in 0.5 seconds.

Comparison with PFE (Prototype Feasibility Estimator, built by BBN based on the Transportation Feasibility Estimator system): On the MEDCOM-SITUATION from the CPE (Common Prototype Environment), KTS is about 5 times faster than PFE and produces a SEA schedule with only 14% of the delay of the PFE schedule. KTS also produces a far more accurate estimate of the planes needed to handle the AIR movements, since PFE is only estimating feasibility whereas KTS produces a detailed schedule.

In our Strategic TPFDD scheduler KTS, we explored issues of speed and embedding KTS into an easy-to-use GUI, complete with ability to edit the data model (TPFDD, resource classes and instances, and port models), to schedule, apply various analysis tools, and to dynamically reschedule. KTS is available from Kestrel via ftp to participants in the PI.

5.4 ITAS – In-Theater Airlift Scheduler

In 1994 we began to develop a scheduler to support PACAF (Pacific Air Force) at Hickham AFB, Honolulu which is tasked with in-theater scheduling of a fleet of 26 C-130 cargo aircraft in the Pacific region. Several variants of a theater scheduler, called ITAS for In-Theater Airlift Scheduler, have been developed to date, and more are planned. The system runs on laptop computers (Apple Powerbook). The interface to ITAS and integration with a commercial database package have been developed by BBN. Users enter information about movement requirements, available resources, port features, etc. and ITAS automatically generates a schedule, displayed in a gantt-like “rainbow” chart. The schedule can also be printed in the form of ATO’s (Air Tasking Orders).

The ITAS schedulers have emphasized flexibility and rich constraint modeling. The version of ITAS installed at PACAF in February 1995 simultaneously schedules the following types of resources:

1. aircraft
2. aircrews and their duty day cycles
3. ground crews for unloading
4. parking space at ports

each of which may have a variety of attendant constraints and problem features.

6 Classification Approach to Algorithm Design

In this section we introduce a new knowledge-based approach to algorithm design. We have been developing it in order to support the incremental application of problem-solving methods to scheduling problems. Our techniques enable us to integrate at a deep semantic level problem-solving methods from Computer Science (e.g. divide-and-conquer, global search), Artificial Intelligence (e.g. heuristic search, constraint propagation, neural nets), and Operations Research (e.g. Simplex, integer programming, network algorithms). Furthermore these techniques have applications far wider than algorithm design, since they apply to the incremental application of *any* kind of knowledge formally represented in a hierarchy.

6.1 Technical Foundations – Theories

A *theory* (i.e. first-order theory presentation) defines a language and constrains the possible meanings of its symbols by axioms and inference rules. Theories can be used to express many kinds of software-related artifacts, including domain models [49], formal requirements [3, 10, 27, 29], programming languages [6, 15, 19], abstract data types and modules [10, 14, 17], and abstract algorithms [43]. There has been much work on operations for constructing larger theories from smaller theories [3, 7, 30].

A *theory morphism* translates the language of one theory into the language of another theory in a way that preserves theorems. Theory morphisms underlie several aspects of software development, including specification refinement and datatype implementation [5, 27, 30, 51], the binding of parameters in parameterized theories [9, 15], algorithm design [22, 43, 52], and data structure design [32]. There has been work on techniques for composing implementations in a way that reflects the structure of the source specification [3, 30]; however these composition techniques leave open the problem of constructing primitive morphisms.

Theories together with their morphisms define a category.

6.2 Refinement Hierarchy and the Ladder Construction

Abstract programming knowledge can be represented by theories. For example, we showed how to represent divide-and-conquer [33], global search (binary search, backtrack, branch-and-bound) [34], and local search (hillclimbing) [22] as theories. The same approach can be applied to data structures [5], architectures [16], and graphical displays (e.g. Gantt charts).

A collection of problem-solving methods can be organized into a refinement hierarchy using theory morphisms as the refinement arrow [43]. See Figure 6.2. The question emerges of how to access and apply knowledge in such a hierarchy. The answer is illustrated in the “ladder construction” diagram in Figure 9.

The left-hand side of the ladder is a path in the refinement hierarchy of algorithm theories starting at the root (Problem Theory). $Spec_0$ is a given specification theory of a problem.

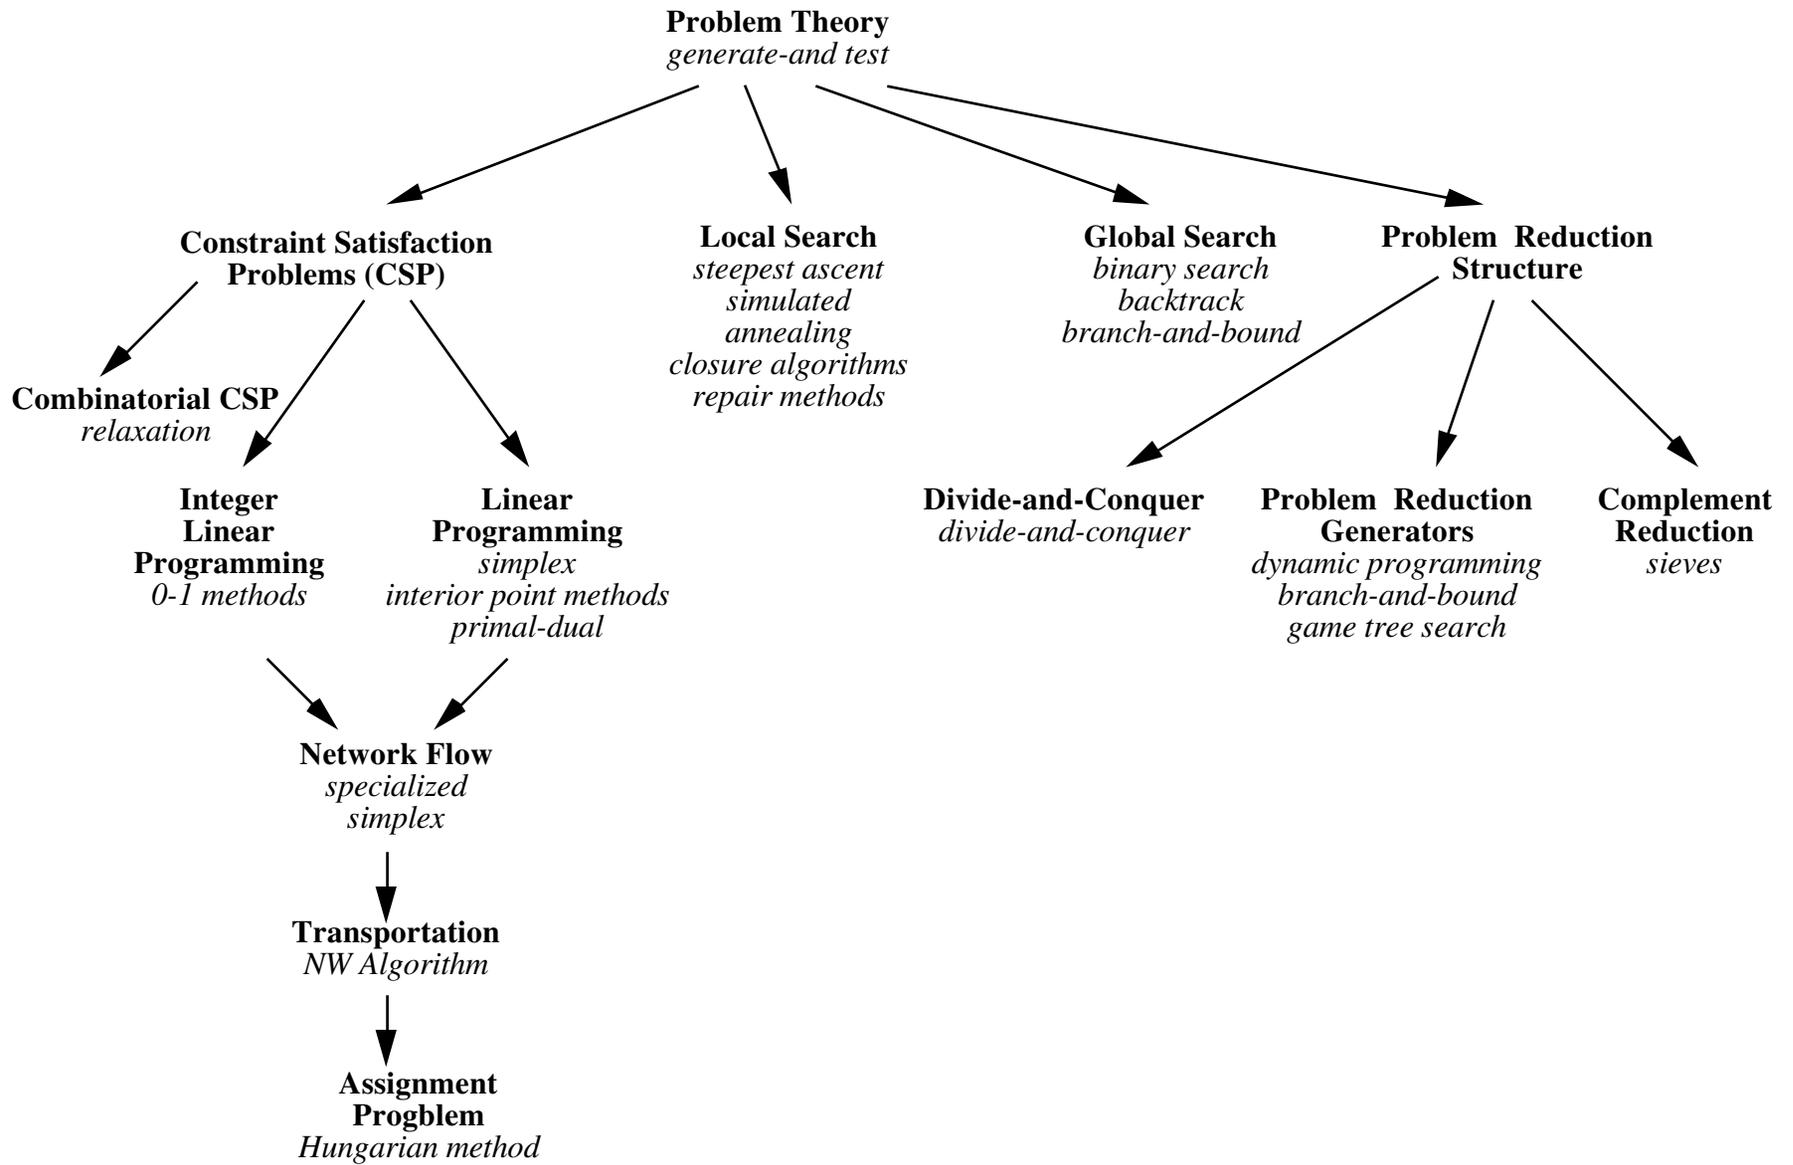


Figure 8: Refinement Hierarchy of Algorithms

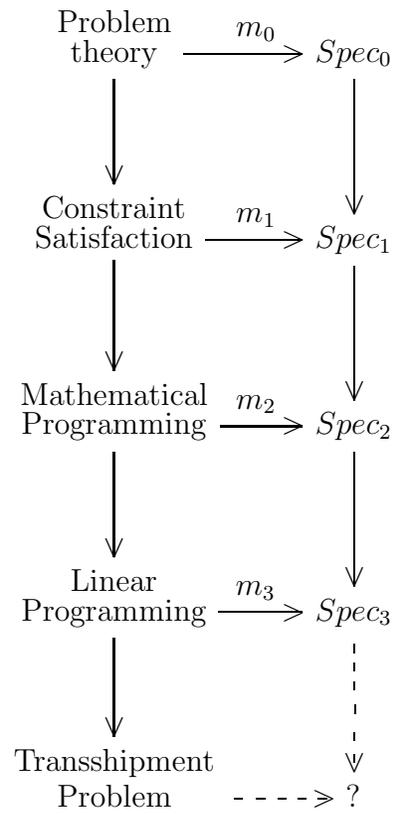


Figure 9: Classification Approach to Design

The ladder is constructed a rung at a time from the top down. The initial arrow (theory morphism) from problem theory to $Spec_0$ is trivial. Subsequent rungs are constructed abstractly as follows:

$$\begin{array}{ccc}
 P_i & \xrightarrow{m_i} & S_i \\
 \downarrow I_i & & \downarrow \\
 P_{i+1} & \xrightarrow{m'_{i+1}} & P_{i+1} \oplus S_i \\
 & \searrow m_{i+1} & \searrow m''_{i+1} \\
 & & S_{i+1}
 \end{array}$$

where $P_{i+1} \oplus S_i$ is the pushout theory (shared union) and S_{i+1} is an extension of S_i determined by constructing the theory morphism m''_{i+1} . The morphism m_{i+1} is determined by composition.

6.3 Constructing Theory Morphisms

Constructing the pushout theory is straightforward. The main issue arising from this ladder construction is how to construct the theory morphism m''_{i+1} from the pushout theory to S_{i+1} (an extension of S_i). We formalized four basic methods for constructing theory morphisms last year, by analyzing the algorithm design tactics in KIDS [39]. Two of the techniques are well-known or obvious. However we identified two new general techniques for constructing theory morphisms: unskolemization and connections between theories. Roughly put, *unskolemization* works in the following way. A theory morphism from theory A to theory B is based on a signature morphism which is a map from the symbols of A to the symbols of B . A theory morphism is signature morphism in which the axioms of A translate to theorems of B . Suppose that during a design process we have somehow managed to construct a partial signature morphism – only some of the symbols of A have a translation as symbols of B . The question is how to derive a translation of the remaining symbols of A . The unskolemization technique uses the axioms of A and deductive inference to solve for appropriate translations of these symbols. As a simple example, suppose that function symbol f is untranslated and that it is the only untranslated symbol in an axiom $\forall(x)G[f(x)]$ of A . We unskolemize f by replacing its occurrence(s) with a fresh existentially quantified variable: $\forall(x)\exists(z)G[z]$. This unskolemized axiom can now be translated and we can attempt to prove it in theory B . A proof yields a witness for the existential that is a term that depends on x . This term can serve as the translation of f knowing that such a translation preserves the theoremhood of the considered axiom. We may need to verify other axioms involving f to assure the appropriateness of the derived translation.

This technique underlies the problem reduction family of algorithms and tactics in KIDS. For example, in constructing a divide-and-conquer algorithm we need to find translations of decompose, solve, and compose operators. The tactic works by letting the user select a standard decomposition operator from a library (or dually, selecting a standard compose operator) and then using unskolemization on a “soundness axiom” that relates decompose and compose. The unskolemized soundness axiom can then be proved in the given problem theory to yield a specification of the compose (resp. decompose) operator. To be more specific, if we are deriving a divide-and-conquer algorithm for the sorting problem, then we want to construct a theory morphism from divide-and-conquer theory into sorting theory. We might choose a standard decompose operator for input sequences, say *split-a-sequence-in-half*, and the unskolemization technique leads to a derivation of a specification for the usual merge operation as the translation of compose. The result is a mergesort algorithm. Other choices leads to quicksort, selection sorts, and insertion sorts [33].

Sometimes the axioms of a theory are too complex to allow direct application of unskolemization. This situation arises in the theory of global and local search algorithms. We have discovered and developed recently the concept of *connection between theories* which underlies and generalizes our correct but somewhat ad-hoc solution to this problem in the global and local search design tactics. The general result regarding connections between theories is this: Suppose that there is a theory T from which we want to construct a theory morphism into a given application domain theory B . If there is a (preexisting) theory morphism from T to a library theory A and we can construct a connection from A to B , then we immediately have a theory morphism from T to B . So connections between theories are a way to adapt a library T -theory to a new, but related problem.

The concept of connection between theories relies on several ideas. The sorts of a theory are all interpreted as posets (including booleans) and furthermore the set of sorts itself is a poset (under the subsort partial order). A collection of “polarity rules” are used to express (anti-)monotonicity properties of the functions and predicates of a theory. For example, $size(\{x \mid \neg P\})$ is monotonic in $\{x \mid \neg P\}$ but antimonotonic in P ; so if $Q \implies P$ then $size(\{x \mid \neg P\}) \leq size(\{x \mid \neg Q\})$. These polarity rules are used to analyze the axioms of a theory and then to set up various connection conditions between the corresponding operators of theories A and B – these conditions directly generalize the conditions of a homomorphism. Furthermore the polarity analysis is used to direct conversion maps between corresponding sorts of A and B . Given these conditions and conversion maps it can in general be shown that the axioms of A imply the corresponding axioms of B , thus establishing the theory morphism.

We have prototyped this classification approach and have tested it on some simple problems. Steve Westfold built a graphical interface to the refinement hierarchy that allows graphical navigation of it and incremental application. Jim McDonald developed a simple Theory Interpretation Construction Interface that supports the development of views (theory interpretations or morphisms). It shows source and target theory presentations and the current (possibly partial) view between them. Users have several tools to support the completion of a view, including typing in translations for various source theory symbols and using “unskolemization” (one of the four basic methods mentioned above). We demonstrated the use of this system to develop a view from divide-and-conquer theory into a simple problem

theory. A more complete implementation of these techniques is underway in the Specware system at Kestrel [50].

7 Concluding Remarks

Our original conception of the scheduling effort has evolved in significant ways. Our 1991 demonstration system was based on use of a general-purpose object base manager and the compilation of declarative constraints into object base demons. We also used a Simplex code to check feasibility of start-times in a generated schedule. The results were somewhat disappointing in that for the CDART problem we obtained from CMU, our first code couldn't solve it running overnight, and our second code could only solve most of it using several minutes time. The derived scheduler described in this paper finds a complete feasible solution to the same problem in less than one second.

Since speed is of the essence during the scheduling process and the object base and Simplex algorithm are problem-independent, it seemed wise to exploit our transformational techniques to try to derive codes that are problem-specific and highly efficient. Rather than compile constraints onto an active database, we now derive pruning mechanisms and constraint propagation code operating on problem-specific data structures. Rather than use a Simplex algorithm for finding feasible start-times, the constraint propagation code maintains feasible start-times throughout the scheduling process. The advantage of our approach is the ability to expose problem structure and exploit it by transformationally deriving efficient problem-specific code.

References

- [1] ALLEN, J. F. Maintaining knowledge about temporal intervals. *Communications of the ACM* 26, 11 (November 1983), 832–843.
- [2] APPLGATE, D., AND COOK, W. A computational study of the job-shop scheduling problem. *ORSA Journal on Computing* 3, 2 (Spring 1991 1991), 149–156.
- [3] ASTESIANO, E., AND WIRSING, M. An introduction to ASL. In *Program Specification and Transformation*, L. Meertens, Ed. North-Holland, Amsterdam, 1987, pp. 343–365.
- [4] BIEFELD, E., AND COOPER, L. Operations mission planner. Tech. Rep. JPL 90-16, Jet Propulsion Laboratory, March 1990.
- [5] BLAINE, L., AND GOLDBERG, A. DTRE – a semi-automatic transformation system. In *Constructing Programs from Specifications*, B. Möller, Ed. North-Holland, Amsterdam, 1991, pp. 165–204.
- [6] BROY, M., WIRSING, M., AND PEPPER, P. On the algebraic definition of programming languages. *ACM Transactions on Programming Languages and Systems* 9, 1 (January 1987), 54–99.

- [7] BURSTALL, R. M., AND GOGUEN, J. A. Putting theories together to make specifications. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence* (Cambridge, MA, August 22–25, 1977), IJCAI, pp. 1045–1058.
- [8] CAI, J., AND PAIGE, R. Program derivation by fixed point computation. *Science of Computer Programming* 11 (1989), 197–261.
- [9] EHRIG, H., KREOWSKI, H. J., THATCHER, J., WAGNER, E., AND WRIGHT, J. Parameter passing in algebraic specification languages. In *Proceedings, Workshop on Program Specification* (Aarhus, Denmark, Aug. 1981), vol. 134, pp. 322–369.
- [10] EHRIG, H., AND MAHR, B. *Fundamentals of Algebraic Specification, vol. 2: Module Specifications and Constraints*. Springer-Verlag, Berlin, 1990.
- [11] ET AL., J. S. *A Review of Strategic Mobility Models and Analysis*. Rand Corporation, Santa Monica, CA, 1991.
- [12] FOX, M. S., SADEH, N., AND BAYKAN, C. Constrained heuristic search. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence* (Detroit, MI, August 20–25, 1989), pp. 309–315.
- [13] FOX, M. S., AND SMITH, S. F. ISIS – a knowledge-based system for factory scheduling. *Expert Systems* 1, 1 (July 1984), 25–49.
- [14] GOGUEN, J. A., THATCHER, J. W., AND WAGNER, E. G. An initial algebra approach to the specification, correctness and implementation of abstract data types. In *Current Trends in Programming Methodology, Vol. 4: Data Structuring*, R. Yeh, Ed. Prentice-Hall, Englewood Cliffs, NJ, 1978.
- [15] GOGUEN, J. A., AND WINKLER, T. Introducing OBJ3. Tech. Rep. SRI-CSL-88-09, SRI International, Menlo Park, California, 1988.
- [16] GRAVES, H. Lockheed environment for automatic programming. Tech. rep., Lockheed Palo Alto Research Center, Palo Alto, CA, 1990.
- [17] GUTTAG, J. V., AND HORNING, J. J. The algebraic specification of abstract data types. *Acta Inf.* 10 (1978), 27–52.
- [18] HENTENRYCK, P. V. *Constraint Satisfaction in Logic Programming*. Massachusetts Institute of Technology, Cambridge, MA, 1989.
- [19] HOARE, C. Varieties of programming languages. Tech. rep., Programming Research Group, University of Oxford, Oxford, UK, 1989.
- [20] JÜLLIG, R., AND SRINIVAS, Y. V. Diagrams for software synthesis. Tech. Rep. KES.U.93.2, Kestrel Institute, March 1993. To appear in: *Proceedings of the 8th Knowledge-Based Software Engineering Conference*, Chicago, IL, September 20–23, 1993.

- [21] LADKIN, P. Specification of time dependencies and synthesis of concurrent processes. In *9th International Conference on Software Engineering* (Monterey, CA, March 30–April 2, 1987), pp. 106–116. Technical Report KES.U.87.1, Kestrel Institute, March 1987.
- [22] LOWRY, M. R. Algorithm synthesis through problem reformulation. In *Proceedings of the 1987 National Conference on Artificial Intelligence* (Seattle, WA, July 13–17, 1987).
- [23] LOWRY, M. R. Automating the design of local search algorithms. In *Automating Software Design*, M. Lowry and R. McCartney, Eds. AAAI/MIT Press, Menlo Park, 1991, pp. 515–546.
- [24] LUENBERGER, D. G. *Linear and Nonlinear Programming*. Addison-Wesley Publishing Company, Inc., Reading, MA, 1989.
- [25] MINTON, S., JOHNSON, M., PHILIPS, A. B., AND LAIRD, P. Solving large-scale constraint satisfaction and scheduling problems using a heuristic repair method. In *Proceedings of the Eighth National Conference on Artificial Intelligence* (1990), pp. 290–295.
- [26] NEMHAUSER, G. L., AND WOLSEY, L. A. *Integer and Combinatorial Optimization*. John Wiley & Sons, Inc., New York, 1988.
- [27] PARTSCH, H. *Specification and Transformation of Programs: A Formal Approach to Software Development*. Springer-Verlag, New York, 1990.
- [28] SADEH, N. Look-ahead techniques for micro-opportunistic job shop scheduling. Tech. Rep. CMU-CS-91-102, Carenegie-Mellon University, March 1991.
- [29] SANNELLA, D., AND TARLECKI, A. Extended ML: An institution-independent framework for formal program development. In *Category Theory and Computer Programming, LNCS 240* (1985), pp. 364–389.
- [30] SANNELLA, D., AND TARLECKI, A. Toward formal development of programs from algebraic specifications: Implementations revisited. *Acta Informatica* 25, 3 (1988), 233–281.
- [31] SELMAN, B., LEVESQUE, H., AND MITCHELL, D. A new method for solving hard satisfiability problems. In *Proceedings of the Tenth National Conference on Artificial Intelligence* (1992), pp. 440–446.
- [32] SMITH, D. R. Data structure design. in preparation.
- [33] SMITH, D. R. Top-down synthesis of divide-and-conquer algorithms. *Artificial Intelligence* 27, 1 (September 1985), 43–96. (Reprinted in *Readings in Artificial Intelligence and Software Engineering*, C. Rich and R. Waters, Eds., Los Altos, CA, Morgan Kaufmann, 1986.).
- [34] SMITH, D. R. Structure and design of global search algorithms. Tech. Rep. KES.U.87.12, Kestrel Institute, November 1987.

- [35] SMITH, D. R. KIDS: A knowledge-based software development system. In *Automating Software Design*, M. Lowry and R. McCartney, Eds. MIT Press, Menlo Park, 1991, pp. 483–514.
- [36] SMITH, D. R. Structure and design of problem reduction generators. In *Constructing Programs from Specifications*, B. Möller, Ed. North-Holland, Amsterdam, 1991, pp. 91–124.
- [37] SMITH, D. R. Transformational approach to scheduling. Tech. Rep. KES.U.92.2, Kestrel Institute, November 1992.
- [38] SMITH, D. R. Classification approach to design. Tech. Rep. KES.U.93.4, Kestrel Institute, 1993.
- [39] SMITH, D. R. Constructing specification morphisms. *Journal of Symbolic Computation, Special Issue on Automatic Programming 15*, 5-6 (May-June 1993), 571–606.
- [40] SMITH, D. R. Toward the synthesis of constraint propagation algorithms. In *Logic Program Synthesis and Transformation (LOPSTR '93)*, Y. DeVille, Ed. Springer-Verlag, 1993, pp. 1–9.
- [41] SMITH, D. R. Synthesis of high-performance transportation schedulers. Tech. Rep. KES.U.95.6, Kestrel Institute, March 1995.
- [42] SMITH, D. R. KIDS – a semi-automatic program development system. *IEEE Transactions on Software Engineering Special Issue on Formal Methods in Software Engineering 16*, 9 (September 1990), 1024–1043.
- [43] SMITH, D. R., AND LOWRY, M. R. Algorithm theories and design tactics. *Science of Computer Programming 14*, 2-3 (October 1990), 305–321.
- [44] SMITH, D. R., AND PARRA, E. A. Transformational approach to transportation scheduling. In *Proceedings of the Eighth Knowledge-Based Software Engineering Conference* (Chicago, IL, September 1993), pp. 60–68.
- [45] SMITH, D. R., AND PARRA, E. A. Transformational approach to transportation scheduling. In *ARPA/RL Knowledge-Based Planning and Scheduling Initiative: Workshop Proceedings* (Tucson, AZ, February 1994), pp. 205–216.
- [46] SMITH, D. R., AND WESTFOLD, S. J. Synthesis of constraint algorithms. In *Principles and Practice of Constraint Programming*, V. Saraswat and P. V. Hentenryck, Eds. The MIT Press, Cambridge, MA, 1995.
- [47] SMITH, S. F. The OPIS framework for modeling manufacturing systems. Tech. Rep. CMU-RI-TR-89-30, The Robotics Institute, Carnegie-Mellon University, December 1989.
- [48] SMITH, S. F., FOX, M. S., AND OW, P. S. Constructing and maintaining detailed production plans: Investigations into the development of knowledge-based factory scheduling systems. *AI Magazine 7*, 4 (Fall 1986), 45–61.

- [49] SRINIVAS, Y. V. Algebraic specification for domains. In *Domain Analysis: Acquisition of Reusable Information for Software Construction*, R. Prieto-Diaz and G. Arango, Eds. IEEE Computer Society Press, Los Alamitos, CA, 1991, pp. 90–124.
- [50] SRINIVAS, Y. V., AND JÜLLIG, R. Specware:tm formal support for composing software. Tech. Rep. KES.U.94.5, Kestrel Institute, December 1994. To appear in Proceedings of the Conference on Mathematics of Program Construction, Kloster Irsee, Germany, July 1995.
- [51] TURSKI, W. M., AND MAIBAUM, T. E. *The Specification of Computer Programs*. Addison-Wesley, Wokingham, England, 1987.
- [52] VELOSO, P. A. Problem solving by interpretation of theories. In *Contemporary Mathematics*, vol. 69. American Mathematical Society, Providence, Rhode Island, 1988, pp. 241–250.
- [53] ZWEBEN, M., DEALE, M., AND GARGAN, R. Anytime rescheduling. In *Proceedings of the Workshop on Innovative Approaches to Planning, Scheduling and Control* (San Diego, CA, November 5–8, 1990), DARPA, pp. 215–219.

A Relation-based Domain Theory for Transportation Scheduling

The following text is the theory of transportation scheduling within which we derived our scheduler. It is expressed in a slight extension of the Refine language. Since Refine uses fairly standard notations for set-theoretic datatypes and first-order logic, most of this text should be generally comprehensible. For more details on domain theories and notation see [42].

This problem theory is based on a schedule as a simple set of reservations:

$$\textit{schedule} : \textit{set}(\textit{movement-record} \times \textit{resource} \times \textit{time})$$

```
THEORY TRANSPORTATION-SCHEDULING
```

```
THEORY-IMPORTS {}
```

```
THEORY-TYPES
```

```
type movement-type = symbol
```

```
type time = integer
```

```
type port = symbol
```

```
type movement-record =
```

```
  tuple(move-type : movement-type,  
        quantity : integer,  
        release-date : time,  
        due-date : time,  
        poe : port,  
        pod : port,  
        mode : symbol,  
        stons : integer,  
        mtons : integer)
```

```
type reservation =
```

```
  tuple(move-rec : movement-record,  
        asset : resource,  
        start-time : time)
```

```
type resource =
```

```
  tuple(pax-capacity : integer,  
        cargo-capacity : integer,  
        class : symbol,
```

```
travel-rate : integer)
```

```
%-----  
THEORY-OPERATIONS
```

```
function DURATION ( res : reservation | true )  
  = distance(POE(res), POD(res)) / travel-rate(asset(res))
```

```
function TRIP-MOVE-RECS ( sched : set(reservation), res : reservation  
  | res in sched )  
  = { move-rec(r) | (r : reservation)  
    r in sched  
    & asset(r) = asset(res)  
    & start-time(r) = start-time(res) }
```

```
function PAX-MOVE-RECS ( sched : set(reservation), res : reservation  
  | res in sched )  
  = filter( lambda(mvr) move-type(mvr) = 'PAX,  
    trip-move-recs(sched, res))
```

```
function CARGO-MOVE-RECS ( sched : set(reservation), res : reservation  
  | res in sched )  
  = filter( lambda(mvr) move-type(mvr) ~= 'PAX,  
    trip-move-recs(sched, res))
```

```
constant cargo-mode-table : map(tuple(movement-type, mode-type), set(resource-class))  
= { |  
  <'pax, 'air>          -> {'C5, 'C141, 'KC10},  
  <'bulk-movement, 'air> -> {'C5, 'C141, 'KC10},  
  <'ovr-movement, 'air> -> {'C5, 'C141},  
  <'out-movement, 'air> -> {'C5},  
  <'pol, 'air>          -> {},  
  <'pax, 'sea>          -> {'Breakbulk, 'Containership, 'Lash, 'Ro-ro, 'Sea-berge},  
  <'bulk-movement, 'sea> -> {'Breakbulk, 'Containership, 'Lash, 'Ro-ro, 'Sea-berge},  
  <'ovr-movement, 'sea> -> {'Breakbulk, 'Containership, 'Lash, 'Ro-ro, 'Sea-berge},  
  <'out-movement, 'sea> -> {'Breakbulk, 'Containership, 'Lash, 'Ro-ro, 'Sea-berge},  
  <'pol, 'sea>          -> {'Large-tanker, 'Medium-tanker, 'Small-tanker}  
| }
```

```
function Consistent-POE  
  ( sched : set(reservation) )  
  = fa (res1 : reservation, res2 : reservation)  
    ( res1 in sched & res2 in sched  
      & asset(res1) = asset(res2)
```

```

    & start-time(res1) = start-time(res2)
    =>
    POE(move-rec(res1)) = POE(move-rec(res2)) )

```

```

function Consistent-POD
( sched : set(reservation) )
= fa (res1 : reservation, res2 : reservation)
    ( res1 in sched & res2 in sched
      & asset(res1) = asset(res2)
      & start-time(res1) = start-time(res2)
      =>
      POD(move-rec(res1)) = POD(move-rec(res2)) )

```

```

function Consistent-Release-Times
(sched : set(reservation) )
= fa (res : reservation)
    ( res in sched
      =>
      release-date(move-rec(res)) <= start-time(res) )

```

```

function Consistent-Due-Times
(sched : set(reservation) )
= fa (res : reservation, mvr : movement-record)
    ( res in sched
      =>
      start-time(res) + duration(res) <= due-date(move-rec(res)) )

```

```

function Consistent-Movement-Type-and-Resource
(sched : set(reservation) )
= fa (res : reservation, mvr : movement-record)
    ( res in sched & mvr = move-rec(res)
      =>
      class(asset(res)) in cargo-mode-table(movetype(mvr),mode(mvr)) )

```

```

function Consistent-Trip-Separation
(sched : set(reservation) )
= fa (res1 : reservation, res2 : reservation)
    ( res1 in sched & res2 in sched
      & res1 ~= res2
      & asset(res1) = asset(res2)
      =>
      start-time(res1) = start-time(res2)
      or start-time(res1) + 2 * duration(res1) <= start-time(res2)
      or start-time(res2) + 2 * duration(res2) <= start-time(res1) )

```

```

function Consistent-Pax-Resource-Capacity

```

```

(sched : set(reservation) )
= fa (res : reservation)
  ( res in sched
    =>
    reduce(+, image(quantity, pax-move-recs(sched, res)))
    <= pax-capacity(asset(res)) )

function Consistent-Cargo-Resource-Capacity
(sched : set(reservation) )
= fa (res : reservation)
  ( res in sched
    =>
    reduce(+, image(quantity, cargo-move-recs(sched, res)))
    <= cargo-capacity(asset(res)) )

function Available-Resources-Used
(assets : set(resource), sched : set(reservation) )
= fa (res : reservation)
  ( res in sched
    =>
    resource(res) in assets )

function All-mvrs-Scheduled
(mvrs : set(movement-record), sched : set(reservation))
= fa (mvr : movement-record)
  ( mvr in mvrs
    =>
    ex!(res : reservation)
    (res in sched & mvr = move-rec(res)) )

function SCHEDULE-MOVEMENT-RECORDS
(mvrs : seq(movement-record),
 assets : set(resource)
 | true)
returns ( sched : set(reservation)) |

  Consistent-POE(sched)
  & Consistent-POD(sched)
  & Consistent-Release-Times(sched)
  & Consistent-Due-Times(sched)
  & Consistent-Movement-Type-and-Resource(sched)
  & Consistent-Trip-Separation(sched)
  & Consistent-Pax-Resource-Capacity(sched)
  & Consistent-Cargo-Resource-Capacity(sched)
  & Available-Resources-Used(assets, sched)

```

```
    & All-mvrs-Scheduled(mvrs, sched)
  )
```

```
%-----
```

```
THEORY-LAWS
```

```
assert extension-of-movement-type
```

```
  fa(mt : movement-type)
```

```
    mt in {'pax, 'bulk-movement, 'ovr-movement, 'out-movement, 'pol}
```

```
...
```

```
end-theory
```

B Map-based Domain Theory for Transportation Scheduling

The following text is the theory of transportation scheduling within which we derived our scheduler. It is expressed in a slight extension of the Refine language. Since Refine uses fairly standard notations for set-theoretic datatypes and first-order logic, most of this text should be generally comprehensible. For more details on domain theories and notation see [42].

This theory represents the bulk of the work needed to derive the scheduler. Subsequent variations and elaborations should be much easier due to reuse of the concepts and laws stated here.

```
THEORY TRANSPORTATION-SCHEDULING-CP
```

```
THEORY-IMPORTS {SEQUENCES-AS-MAPS, TS-RESOURCES}
```

```
THEORY-TYPES
```

```
type RESOURCE =
```

```
  tuple(res-name : symbol,  
        pax-capacity : integer,  
        cargo-capacity : integer,  
        res-class : resource-class,  
        travel-rate : integer)
```

```
type MOVEMENT-RECORD =
```

```
  tuple(move-type : movement-type,  
        quantity : integer,  
        release-date : time,  
        due-date : time,  
        poe : port,  
        pod : port,  
        distance : integer,  
        mode : symbol)
```

```
type resource-name = symbol
```

```
type TRIP =
```

```
  tuple(poe : port,  
        pod : port,  
        start-time : time,  
        trip-duration : integer,  
        manifest : seq(movement-record))
```

```

    )

type P-TRIP =
  tuple(poe : port,
        pod : port,
        earliest-start-time : time,
        latest-start-time : time,
        p-trip-duration : integer,
        manifest : seq(movement-record)
        )

type SCHEDULE = map( resource-name, seq(trip) )
type P-SCHEDULE = map( resource-name, seq(p-trip) )

THEORY-OPERATIONS

function EMPTY-INVERSE-MAP
  (dom : seq(resource-name))
  : p-schedule
  = {| x -> [] | (x : resource-name) x in dom |}

function INVERSE-MAP-APPEND
  (ps : p-schedule,
   r : resource-name,
   trp : p-trip)
  : p-schedule
  = {| x -> ps(x) ++ [trp | () x = r] | (x : resource-name) x in domain(ps) |}

function INVERSE-MAP-WITH!
  (ps : p-schedule,
   r : resource-name,
   trp : p-trip,
   mvr : movement-record
   | trp in ps(r))
  : p-schedule
  = {| r-var -> [<emb, deb,
                (if r = r-var & trp = trp-var
                  then max(est, mvr.release-date) else est),
                (if r = r-var & trp = trp-var
                  then min(lst,
                          mvr.due-date - duration(mvr.distance, *resource-map*(r-var)
                          else lst),
                dur,
                (if r = r-var & trp = trp-var then prepend(mani, mvr)
                  else mani)>
   | (emb : port, deb : port,

```

```

        est : time, lst : time,
        trp-var : p-trip, dur : integer,
        mani : seq(movement-record))
trp-var in ps(r-var)
    & emb = trp-var.poe
    & deb = trp-var.pod
    & est = trp-var.earliest-start-time
    & lst = trp-var.latest-start-time
    & dur = trp-var.p-trip-duration
    & mani = trp-var.manifest ]
| (r-var : resource-name) r-var in domain(ps) |}

% other input conditions needed (e.g. matching poe, pod, est, lst per trip)
% also, need to combine the est and lst intervals
function INVERSE-MAP-UNION!
    (ps : p-schedule,
     qs : p-schedule
     | domain(ps) = domain(qs))
    : p-schedule
    = {| x -> [<emb, deb, est, lst, dur,
              ( ps(x)(trp-var).manifest ++ qs(x)(trp-var).manifest)>
              | (emb : port, deb : port,
                est : time, lst : time, dur : integer,
                trp-var : integer)
              trp-var in [1 .. size(ps(x))]
              & emb = ps(x)(trp-var).poe
              & deb = ps(x)(trp-var).pod
              & est = ps(x)(trp-var).earliest-start-time
              & lst = ps(x)(trp-var).latest-start-time
              & dur = ps(x)(trp-var).p-trip-duration]
              | (x:resource-name) x in domain(ps) |}]

% the extracted schedule uses the earliest-start-time as the start-time
% for each trip. This of course is arbitrary.
function EXTRACT-A-SCHEDULE
    (p-sched : p-schedule) : schedule
    = {| r -> [ <p-trp.poe,
              p-trp.pod,
              p-trp.earliest-start-time,
              p-trp.p-trip-duration,
              p-trp.manifest>
              | (p-trp : p-trip) p-trp in p-sched(r) ]
          | (r : resource-name) r in domain(p-sched)
          |}]

```

```

function INVERSE-MAP?
  (assets : seq(resource-name),
   ps : p-schedule
   | domain(ps) = assets): boolean
= fa(r1 : resource-name, r2 : resource-name, trp1 : p-trip, trp2 : p-trip)
  (r1 in assets & r2 in assets & r1 ~= r2
   & trp1 in ps(r1) & trp2 in ps(r2)
   => disjoint(seq-to-set(trp1.manifest), seq-to-set(trp2.manifest)))
&
fa(r : resource-name, trp1 : p-trip, trp2 : p-trip)
  (r in assets
   & trp1 in ps(r) & trp2 in ps(r) & trp1 ~= trp2
   => disjoint(seq-to-set(trp1.manifest), seq-to-set(trp2.manifest)))

function INVERSE-MAP-SATISFIES
  (sched : schedule,
   p-sched : p-schedule) : boolean
= ( domain(p-sched) = domain(sched)
   & fa(r : resource-name, trp : integer)
     ( r in domain(p-sched)
      & trp in [1 .. size(p-sched(r))]
      =>
        ( p-sched(r)(trp).manifest subset sched(r)(trp).manifest
          & p-sched(r)(trp).poe = sched(r)(trp).poe
          & p-sched(r)(trp).pod = sched(r)(trp).pod
          & p-sched(r)(trp).p-trip-duration = sched(r)(trp).trip-duration
          & p-sched(r)(trp).earliest-start-time <= sched(r)(trp).start-time
          & sched(r)(trp).start-time <= p-sched(r)(trp).latest-start-time
          )))

function seq-extends?(s1: seq(-alpha), s2: seq(-alpha)): boolean =
  [.., $s2] = s1

% test if q-sched is a refinement of p-sched
function LESS-CONSTRAINED-THAN
  (p-sched : p-schedule,
   q-sched : p-schedule) : boolean
= ( domain(p-sched) subset domain(q-sched)
   & fa(r : resource-name)
     ( r in domain(p-sched)
      => size(p-sched(r)) <= size(q-sched(r))
      & fa(trp : integer)
        (trp in [1 .. size(p-sched(r))]
         =>
           ( seq-extends?(q-sched(r)(trp).manifest,
                          p-sched(r)(trp).manifest)
            )))

```

```

        & p-sched(r)(trp).earliest-start-time
          <= q-sched(r)(trp).earliest-start-time
        & q-sched(r)(trp).latest-start-time
          <= p-sched(r)(trp).latest-start-time
      )))
  )

```

%% Properties of Schedules %%%

```

function CONSISTENT-POE
  ( sched : schedule ) : boolean
= fa ( rsrc : resource-name, trp : integer, mvr : movement-record)
    ( rsrc in domain(sched) & trp in [1 .. size(sched(rsrc))]
      & mvr in sched(rsrc)(trp).manifest
    =>
      mvr.poe = sched(rsrc)(trp).poe )

```

```

function CONSISTENT-POD
  ( sched : schedule ) : boolean
= fa ( rsrc : resource-name, trp : integer, mvr : movement-record)
    ( rsrc in domain(sched) & trp in [1 .. size(sched(rsrc))]
      & mvr in sched(rsrc)(trp).manifest
    =>
      mvr.pod = sched(rsrc)(trp).pod )

```

```

function CONSISTENT-RELEASE-TIMES
  ( sched : schedule ) : boolean
= fa ( rsrc : resource-name, trp : integer, mvr : movement-record)
    ( rsrc in domain(sched)
      & trp in [1 .. size(sched(rsrc))]
      & mvr in sched(rsrc)(trp).manifest
    =>
      mvr.release-date <= sched(rsrc)(trp).start-time )

```

```

function CONSISTENT-DUE-TIMES
  ( sched : schedule ) : boolean
= fa ( rsrc : resource-name, trp : integer, mvr : movement-record)
    ( rsrc in domain(sched)
      & trp in [1 .. size(sched(rsrc))]
      & mvr in sched(rsrc)(trp).manifest
    =>
      sched(rsrc)(trp).start-time
        <= (mvr.due-date - sched(rsrc)(trp).trip-duration) )

```

```

function CONSISTENT-TRIP-SEPARATION
(sched : schedule) : boolean
= fa (rsrc : resource-name, trp : integer)
  ( rsrc in domain(sched)
    & trp in [1 .. size(sched(rsrc)) - 1]
    =>
      sched(rsrc)(trp).start-time + 2 * sched(rsrc)(trp).trip-duration
      <= sched(rsrc)(trp + 1).start-time )

```

```

function pax-resource-used-in-mvrs
(mvrs : seq(movement-record)) : integer
= reduce(+, image(lambda(mvr: movement-record) mvr.quantity,
                  filter( lambda(mvr) mvr.move-type = 'pax, mvrs)))

```

```

function CONSISTENT-PAX-RESOURCE-CAPACITY
(sched : schedule) : boolean
= fa (rsrc : resource-name, trp : integer)
  ( rsrc in domain(sched) & trp in [1 .. size(sched(rsrc))])
  =>
    pax-resource-used-in-mvrs(sched(rsrc)(trp).manifest)
    <= *resource-map*(rsrc).pax-capacity )

```

```

function cargo-resource-used-in-mvrs
(mvrs : seq(movement-record)) : integer
= reduce(+, image(lambda(mvr: movement-record) mvr.quantity,
                  filter( lambda(mvr) mvr.move-type ~= 'pax, mvrs)))

```

```

function CONSISTENT-CARGO-RESOURCE-CAPACITY
(sched : schedule) : boolean
= fa (rsrc : resource-name, trp : integer)
  ( rsrc in domain(sched) & trp in [1 .. size(sched(rsrc))])
  =>
    cargo-resource-used-in-mvrs(sched(rsrc)(trp).manifest)
    <= *resource-map*(rsrc).cargo-capacity )

```

```

function CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE
(sched : schedule) : boolean
= fa (rsrc : resource-name, trp : integer, mvr : movement-record)
  ( rsrc in domain(sched) & trp in [1 .. size(sched(rsrc))]
    & mvr in sched(rsrc)(trp).manifest
    =>
      *resource-map*(rsrc).res-class
      in cargo-mode-table(mvr.move-type, mvr.mode) )

```

```

function AVAILABLE-RESOURCES-USED
  (assets : seq(resource-name), sched : schedule) : boolean
  = domain(sched) subset seq-to-set(assets)

% need inverse map operations to help express
% that the domain of the inverse map is exactly mvrs

function SCHEDULED-MVRS
  (sched : schedule) : set(movement-record)
  = seq-to-set(reduce(concat,
                    image(lambda(x: trip) x.manifest,
                          reduce(concat, range(sched))))))

function TS
  (mvrs : seq(movement-record),
   assets : seq(resource-name)
   | true)
  returns ( sched : map( resource-name, seq(trip) ) |

          Consistent-POE(sched)
          & Consistent-POD(sched)

          & Consistent-Release-Times(sched)
          & Consistent-Due-Times(sched)
          & Consistent-Trip-Separation(sched)

          & Consistent-Pax-Resource-Capacity(sched)
          & Consistent-Cargo-Resource-Capacity(sched)

          & Consistent-Movement-Type-and-Resource(sched)
          & Available-Resources-Used(assets, sched)

          & Scheduled-mvrs(sched) = seq-to-set(mvrs)
          )

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Properties of P-Schedules %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

function P-CONSISTENT-POE ( p-sched : p-schedule) : boolean
  = fa ( rsrc : resource-name, trp : integer, mvr : movement-record)
      ( rsrc in domain(p-sched) & trp in [1 .. size(p-sched(rsrc))]
        & mvr in p-sched(rsrc)(trp).manifest
        =>
          mvr.poe = p-sched(rsrc)(trp).poe )

function P-CONSISTENT-POD

```

```

( p-sched : p-schedule ) : boolean
= fa ( rsrc : resource-name, trp : integer, mvr : movement-record)
  ( rsrc in domain(p-sched) & trp in [1 .. size(p-sched(rsrc))]
    & mvr in p-sched(rsrc)(trp).manifest
  =>
  mvr.pod = p-sched(rsrc)(trp).pod )

function P-CONSISTENT-RELEASE-TIMES
(p-sched : p-schedule) : boolean
= fa ( rsrc : resource-name, p-trp : integer, mvr : movement-record)
  ( rsrc in domain(p-sched)
    & p-trp in [1 .. size(p-sched(rsrc))]
    & mvr in p-sched(rsrc)(p-trp).manifest
  =>
  mvr.release-date <= p-sched(rsrc)(p-trp).earliest-start-time )

function P-CONSISTENT-DUE-TIMES
(p-sched : p-schedule) : boolean
= fa ( rsrc : resource-name, p-trp : integer, mvr : movement-record)
  ( rsrc in domain(p-sched)
    & p-trp in [1 .. size(p-sched(rsrc))]
    & mvr in p-sched(rsrc)(p-trp).manifest
  =>
  p-sched(rsrc)(p-trp).latest-start-time
  <= (mvr.due-date - p-sched(rsrc)(p-trp).p-trip-duration))

% the trick is to derive and enforce this invariant

function P-CONSISTENT-TRIP-SEPARATION-EST
(ps :p-schedule) : boolean
= fa ( rsrc : resource-name, trp : integer)
  ( rsrc in domain(ps) & trp in [1 .. size(ps(rsrc)) - 1]
  => ps(rsrc)(trp).earliest-start-time
    + 2 * ps(rsrc)(trp).p-trip-duration
    <= ps(rsrc)(trp + 1).earliest-start-time )

function P-CONSISTENT-TRIP-SEPARATION-LST
(ps :p-schedule) : boolean
= fa ( rsrc : resource-name, trp : integer)
  ( rsrc in domain(ps) & trp in [1 .. size(ps(rsrc)) - 1]
  => ps(rsrc)(trp).latest-start-time
    <= ps(rsrc)(trp + 1).latest-start-time
    - 2 * ps(rsrc)(trp).p-trip-duration)

function PAX-RESOURCE-USED-IN-MVRS

```

```

(mvrs : seq(movement-record)) : integer
= reduce(+, image(lambda(mvr: movement-record) mvr.quantity,
                  filter( lambda(mvr) mvr.move-type = 'Pax, mvrs)))

function P-CONSISTENT-PAX-RESOURCE-CAPACITY
(p-sched : p-schedule) : boolean
= fa (rsrc : resource-name, trp : integer)
  ( rsrc in domain(p-sched) & trp in [1 .. size(p-sched(rsrc))]
    => %p-sched(rsrc)(trp).pax-used
      pax-resource-used-in-mvrs(p-sched(rsrc)(trp).manifest)
      <= *resource-map*(rsrc).pax-capacity )

function P-CONSISTENT-CARGO-RESOURCE-CAPACITY
(p-sched : p-schedule) : boolean
= fa (rsrc : resource-name, trp : integer)
  ( rsrc in domain(p-sched) & trp in [1 .. size(p-sched(rsrc))]
    =>
      % p-sched(rsrc)(trp).cargo-used
      cargo-resource-used-in-mvrs(p-sched(rsrc)(trp).manifest)
      <= *resource-map*(rsrc).Cargo-capacity )

function P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE
(p-sched : p-schedule) : boolean
= fa (rsrc : resource-name, trp : integer, mvr : movement-record)
  ( rsrc in domain(p-sched) & trp in [1 .. size(p-sched(rsrc))]
    & mvr in p-sched(rsrc)(trp).manifest
    =>
      *resource-map*(rsrc).res-class
      in cargo-mode-table(mvr.move-type, mvr.mode) )

function P-AVAILABLE-RESOURCES-USED
(assets : seq(resource-name), p-sched : p-schedule) : boolean
= fa (rsrc : resource-name) ( rsrc in domain(p-sched) => rsrc in assets)

function P-SCHEDULED-MVRS
(p-sched : p-schedule) : set(movement-record)
= seq-to-set(reduce(concat, image(lambda(x: p-trip) x.manifest,
                                  reduce(concat, range(p-sched))))))

function ALL-MVRS-P-SCHEDULED
(mvrs : set(movement-record),
 p-sched : p-schedule) : boolean
= p-scheduled-mvrs(p-sched) = mvrs

function P-CONSISTENT-TIME-INTERVALS
(p-sched : p-schedule) : boolean

```

```

= fa(r : resource-name, trp : p-trip)
  (r in domain(p-sched)
   & trp in p-sched(r)
   => trp.earliest-start-time <= trp.latest-start-time)

```

THEORY-LAWS

```

assert definedness-of-empty-inverse-map
  fa(assets)( defined?(empty-inverse-map(assets)) = defined?(assets))

```

```

assert universal-lower-bound-of-less-constrained-than
  fa(a,b)( less-constrained-than(empty-inverse-map(a), b) = true )

```

```

assert EXTRACT-A-SCHEDULE-PRODUCES-AN-INSTANCE-OF-INVERSE-MAP
  fa(ps)( inverse-map-satisfies(extract-a-schedule(ps),ps) = true )

```

```

assert INVERSE-MAP-SATISFIES-IMPLIES-P-CONSISTENT-TIME-INTERVALS
  fa(sched, ps)( inverse-map-satisfies(sched, ps)
                 => p-consistent-time-intervals(ps))

```

```

assert CONSISTENT-POE-to-P-CONSISTENT-POE
  fa(S,PS)( inverse-map-satisfies(S,PS)
            => (Consistent-POE(S) => P-Consistent-POE(PS)) )

```

```

assert CONSISTENT-POD-to-P-CONSISTENT-POD
  fa(S,PS)( inverse-map-satisfies(S,PS)
            => (Consistent-POD(S) => P-Consistent-POD(PS)) )

```

```

assert CONSISTENT-PAX-RESOURCE-CAPACITY-to-P-CONSISTENT-PAX-RESOURCE-CAPACITY
  fa(S,PS)( inverse-map-satisfies(S,PS)
            => (CONSISTENT-PAX-RESOURCE-CAPACITY(S)
                => P-CONSISTENT-PAX-RESOURCE-CAPACITY(PS)) )

```

```

assert CONSISTENT-CARGO-RESOURCE-CAPACITY-to-P-CONSISTENT-CARGO-RESOURCE-CAPACITY
  fa(S,PS)( inverse-map-satisfies(S,PS)
            => (CONSISTENT-CARGO-RESOURCE-CAPACITY(S)
                => P-CONSISTENT-CARGO-RESOURCE-CAPACITY(PS)) )

```

```

assert CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-to-P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE
  fa(S,PS)( inverse-map-satisfies(S,PS)
            => (CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE(S)
                => P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE(PS)) )

```

```

assert AVAILABLE-RESOURCES-USED-to-P-AVAILABLE-RESOURCES-USED
  fa(assets,S,PS)
    ( inverse-map-satisfies(S,PS)
      => ( AVAILABLE-RESOURCES-USED(assets,S)
          => P-AVAILABLE-RESOURCES-USED(assets,PS)) )

assert SCHEDULED-MVRS-to-P-SCHEDULED-MVRS
  fa(S,PS)( inverse-map-satisfies(S,PS)
            => (SCHEDULED-MVRS(S) => P-SCHEDULED-MVRS(PS)) )

assert CONSISTENT-RELEASE-TIMES-to-P-CONSISTENT-RELEASE-TIMES
  fa(PS)( fa(S)(inverse-map-satisfies(S,PS) => CONSISTENT-RELEASE-TIMES(S))
          => P-CONSISTENT-RELEASE-TIMES(PS))

assert CONSISTENT-DUE-TIMES-to-P-CONSISTENT-DUE-TIMES
  fa(PS)( fa(S)(inverse-map-satisfies(S,PS) => CONSISTENT-DUE-TIMES(S))
          => P-CONSISTENT-DUE-TIMES(PS))

assert CONSISTENT-TRIP-EST-AND-LST-SEPARATION-to-P-CONSISTENT-TRIP-SEPARATION
  fa(PS)( fa(S)(INVERSE-MAP-SATISFIES(S, PS) => CONSISTENT-TRIP-SEPARATION(S))
          =>
            (P-CONSISTENT-TRIP-SEPARATION-EST(PS) & P-CONSISTENT-TRIP-SEPARATION-LST(PS)))

assert CONSISTENT-POE-over-extract
  fa(S,PS)( Consistent-POE(extract-a-schedule(ps)) = P-Consistent-POE(PS) )

assert CONSISTENT-POD-over-extract
  fa(S,PS)(Consistent-POD(Extract-a-schedule(Ps)) = P-Consistent-POD(PS))

assert CONSISTENT-RELEASE-TIMES-over-extract
  fa(S,PS)(P-CONSISTENT-RELEASE-TIMES(PS)
          => CONSISTENT-RELEASE-TIMES(Extract-a-schedule(Ps)))

assert CONSISTENT-DUE-TIMES-over-extract
  fa(S,PS)(P-CONSISTENT-DUE-TIMES(PS) => CONSISTENT-DUE-TIMES(Extract-a-schedule(Ps)) )

assert CONSISTENT-TRIP-SEPARATION-over-extract
  fa(s,ps)(p-consistent-trip-separation-est(ps) & p-consistent-trip-separation-lst(ps)
          => consistent-trip-separation(extract-a-schedule(ps)) )

assert CONSISTENT-PAX-RESOURCE-CAPACITY-over-extract
  fa(S,PS)(CONSISTENT-PAX-RESOURCE-CAPACITY(Extract-a-schedule(Ps))

```

```

= P-CONSISTENT-PAX-RESOURCE-CAPACITY(PS))

assert CONSISTENT-CARGO-RESOURCE-CAPACITY-over-extract
  fa(S,PS) (CONSISTENT-CARGO-RESOURCE-CAPACITY(Extract-a-schedule(Ps))
    = P-CONSISTENT-CARGO-RESOURCE-CAPACITY(PS))

assert CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-over-extract
  fa(S,PS) (CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE(Extract-a-schedule(Ps))
    = P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE(PS))

assert AVAILABLE-RESOURCES-USED-over-extract
  fa(assets,S,PS)
    ( AVAILABLE-RESOURCES-USED(assets,Extract-a-schedule(ps))
      = P-AVAILABLE-RESOURCES-USED(assets,PS) )

assert SCHEDULED-MVRS-over-extract
  fa(S,PS) (SCHEDULED-MVRS(Extract-a-schedule(ps)) = P-SCHEDULED-MVRS(PS))

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Distributive Laws %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
assert DISTRIBUTE-P-CONSISTENT-POE-OVER-EMPTY-INVERSE-MAP
  fa(dom) ( p-consistent-poe(EMPTY-INVERSE-MAP(dom)) = true )

assert DISTRIBUTE-P-CONSISTENT-POE-OVER-INVERSE-MAP-with!
  fa(ps,r,trp,mvr) ( p-consistent-poe(inverse-map-with!(ps,r,trp,mvr))
    = (p-consistent-poe(ps) & trp.poe = mvr.poe))

assert DISTRIBUTE-P-CONSISTENT-POE-OVER-INVERSE-MAP-append
  fa(ps,r,trp) ( p-consistent-poe(inverse-map-append(ps,r,trp))
    = p-consistent-poe(ps))

assert DISTRIBUTE-P-CONSISTENT-POE-OVER-INVERSE-MAP-UNION!
  fa(ps,qs) ( p-consistent-poe(inverse-map-union!(ps,qs))
    = (p-consistent-poe(ps) & p-consistent-poe(qs)) )

assert DISTRIBUTE-P-CONSISTENT-POD-OVER-EMPTY-INVERSE-MAP
  fa(dom) ( P-consistent-pod(EMPTY-INVERSE-MAP(dom)) = true )

assert DISTRIBUTE-P-CONSISTENT-POD-OVER-INVERSE-MAP-with!
  fa(ps,r,trp,mvr) ( P-consistent-pod(inverse-map-with!(ps,r,trp,mvr))
    = (P-consistent-pod(ps) & trp.pod = mvr.pod))

assert DISTRIBUTE-P-CONSISTENT-POD-OVER-INVERSE-MAP-append

```

```

fa(ps,r,trp)( P-consistent-pod(inverse-map-append(ps,r,trp))
              = P-consistent-pod(ps))

assert DISTRIBUTE-P-CONSISTENT-POD-OVER-INVERSE-MAP-UNION!
fa(ps,qs)( P-consistent-pod(inverse-map-union!(ps,qs))
          = (P-consistent-pod(ps) & p-consistent-POD(qs)) )

assert DISTRIBUTE-P-CONSISTENT-PAX-RESOURCE-CAPACITY-OVER-EMPTY-INVERSE-MAP
fa(dom)( P-CONSISTENT-PAX-RESOURCE-CAPACITY(EMPTY-INVERSE-MAP(dom))
        = true )

assert DISTRIBUTE-P-CONSISTENT-PAX-RESOURCE-CAPACITY-OVER-INVERSE-MAP-with!
fa(ps,r,trp: p-trip,mvr)
( P-CONSISTENT-PAX-RESOURCE-CAPACITY(inverse-map-with!(ps,r,trp,mvr))
  = (P-CONSISTENT-PAX-RESOURCE-CAPACITY(ps)
    & (mvr.move-type ~= 'PAX
      or-else pax-resource-used-in-mvrs(trp.manifest) + mvr.quantity
      <= *resource-map*(r).pax-capacity)))

assert DISTRIBUTE-P-CONSISTENT-PAX-RESOURCE-CAPACITY-OVER-INVERSE-MAP-append
fa(ps,r,trp: p-trip)
( P-CONSISTENT-PAX-RESOURCE-CAPACITY(inverse-map-append(ps,r,trp))
  = (P-CONSISTENT-PAX-RESOURCE-CAPACITY(ps)
    & (arb(trp.manifest).move-type ~= 'PAX
      or-else pax-resource-used-in-mvrs(trp.manifest)
      <= *resource-map*(r).pax-capacity)))

assert DISTRIBUTE-P-CONSISTENT-PAX-RESOURCE-CAPACITY-OVER-INVERSE-MAP-UNION!
fa(ps,qs)( P-CONSISTENT-PAX-RESOURCE-CAPACITY(inverse-map-union!(ps,qs))
          => P-CONSISTENT-PAX-RESOURCE-CAPACITY(ps) )

assert DISTRIBUTE-P-CONSISTENT-CARGO-RESOURCE-CAPACITY-OVER-EMPTY-INVERSE-MAP
fa(dom)( P-CONSISTENT-CARGO-RESOURCE-CAPACITY(EMPTY-INVERSE-MAP(dom))
        = true )

assert DISTRIBUTE-P-CONSISTENT-CARGO-RESOURCE-CAPACITY-OVER-INVERSE-MAP-with!
fa(ps,r,trp: p-trip, mvr)
( P-CONSISTENT-CARGO-RESOURCE-CAPACITY(inverse-map-with!(ps,r,trp,mvr))
  = (P-CONSISTENT-CARGO-RESOURCE-CAPACITY(ps)
    & (mvr.move-type = 'PAX
      or-else cargo-resource-used-in-mvrs(trp.manifest) + mvr.quantity

```

```

                                <= *resource-map*(r).cargo-capacity)))

assert DISTRIBUTE-P-CONSISTENT-CARGO-RESOURCE-CAPACITY-OVER-INVERSE-MAP-append
  fa(ps,r,trp: p-trip)
  ( P-CONSISTENT-CARGO-RESOURCE-CAPACITY(inverse-map-append(ps,r,trp))
    = (P-CONSISTENT-CARGO-RESOURCE-CAPACITY(ps)
      & (arb(trp.manifest).move-type = 'PAX
        or-else cargo-resource-used-in-mvrs(trp.manifest)
        <= *resource-map*(r).cargo-capacity)) )

assert DISTRIBUTE-P-CONSISTENT-CARGO-RESOURCE-CAPACITY-OVER-INVERSE-MAP-UNION!
  fa(ps,qs)( P-CONSISTENT-CARGO-RESOURCE-CAPACITY(inverse-map-union!(ps,qs))
    => P-CONSISTENT-CARGO-RESOURCE-CAPACITY(ps) )

assert DISTRIBUTE-P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-OVER-EMPTY-INVERSE-MAP
  fa(dom)( p-consistent-movement-type-and-resource(EMPTY-INVERSE-MAP(dom))
    = true )

assert DISTRIBUTE-P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-OVER-INVERSE-MAP-with!
  fa(ps,r,trp,mvr)
  ( p-consistent-movement-type-and-resource(inverse-map-with!(ps,r,trp,mvr))
    = (p-consistent-movement-type-and-resource(ps)
      & *resource-map*(r).res-class in cargo-mode-table(mvr.move-type, mvr.mode)) )

assert DISTRIBUTE-P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-OVER-INVERSE-MAP-append
  fa(ps,r,trp)
  ( p-consistent-movement-type-and-resource(inverse-map-append(ps,r,trp))
    = (p-consistent-movement-type-and-resource(ps)
      & *resource-map*(r).res-class in cargo-mode-table(arb(trp.manifest).move-type,
        arb(trp.manifest).mode)) )

assert DISTRIBUTE-P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-OVER-INVERSE-MAP-UNION!
  fa(ps,qs)( p-consistent-movement-type-and-resource(inverse-map-union!(ps,qs))
    = (p-consistent-movement-type-and-resource(ps)
      & p-consistent-movement-type-and-resource(qs)) )

assert DISTRIBUTE-P-AVAILABLE-RESOURCES-USED-OVER-EMPTY-INVERSE-MAP
  fa(assets)( P-AVAILABLE-RESOURCES-USED(assets, EMPTY-INVERSE-MAP(assets))
    = true )

assert DISTRIBUTE-P-AVAILABLE-RESOURCES-USED-OVER-INVERSE-MAP-with!
  fa(assets,ps,r,trp,mvr)

```

```

    ( P-AVAILABLE-RESOURCES-USED(assets, inverse-map-with!(ps,r,trp,mvr))
      = (P-AVAILABLE-RESOURCES-USED(assets,ps) & r in assets) )

assert DISTRIBUTE-P-AVAILABLE-RESOURCES-USED-OVER-INVERSE-MAP-append
  fa(assets,ps,r,trp)
    ( P-AVAILABLE-RESOURCES-USED(assets, inverse-map-append(ps,r,trp))
      = (P-AVAILABLE-RESOURCES-USED(assets,ps) & r in assets) )

assert DISTRIBUTE-P-AVAILABLE-RESOURCES-USED-OVER-INVERSE-MAP-UNION!
  fa(assets,ps,qs) ( P-AVAILABLE-RESOURCES-USED(assets, inverse-map-union!(ps,qs))
                    = (P-AVAILABLE-RESOURCES-USED(assets,ps)
                      & P-AVAILABLE-RESOURCES-USED(assets,qs)) )

assert DISTRIBUTE-ALL-MVRS-P-SCHEDULED-OVER-EMPTY-INVERSE-MAP
  fa(assets) ( ALL-MVRS-P-SCHEDULED({}, EMPTY-INVERSE-MAP(assets)) = true )

assert DISTRIBUTE-P-CONSISTENT-TIME-INTERVALS-OVER-EMPTY-INVERSE-MAP
  fa(dom) ( p-consistent-time-intervals(empty-inverse-map(dom)) = true )

assert DISTRIBUTE-P-CONSISTENT-TIME-INTERVALS-OVER-INVERSE-MAP-WITH!
  fa (ps,r,p-trp,mvr)
    ( p-consistent-time-intervals(inverse-map-with!(ps,r,p-trp,mvr))
      = (p-consistent-time-intervals(ps)
        & p-trp.earliest-start-time <= mvr.due-date - p-trp.p-trip-duration
        & mvr.release-date <= p-trp.latest-start-time))

assert DISTRIBUTE-P-CONSISTENT-TIME-INTERVALS-OVER-INVERSE-MAP-APPEND
  fa(ps,r,trp)
    ( p-consistent-time-intervals(inverse-map-append(ps,r,trp))
      = p-consistent-time-intervals(ps))

assert PAX-RESOURCE-USED-IN-MVRS-OVER-SINGLETON-MANIFEST
  fa(e,d,est,lst,dur,mvr)
    ( pax-resource-used-in-mvrs({mvr}) = mvr.quantity)

assert CARGO-RESOURCE-USED-IN-MVRS-OVER-SINGLETON-MANIFEST
  fa(e,d,est,lst,dur,mvr)
    ( cargo-resource-used-in-mvrs({mvr}) = mvr.quantity)

```

```

assert SELECT-MANIFEST-FROM-LITERAL-TRIP
  fa(e,d,s,dur,m) ( <e,d,s,dur,m>.manifest = m )

assert SELECT-MANIFEST-FROM-LITERAL-P-TRIP
  fa(e,d,est,lst,dur,m) ( <e,d,est,dur,lst,m>.manifest = m )

```

```

assert member-seq-range-true
  fa(x) ((x in [1 .. x]) = (x >= 1))

assert member-seq-range-minus-1
  fa(x) ((x in [1 .. x + -1]) = false)

assert member-plus-1-seq-range
  fa(x) ((x + 1 in [1 .. x]) = false)

```

THEORY-RULES

```

function ts-rule-member-seq-range-true ()
  rb-compile-simplification-equality
  member-seq-range-true

```

```

function ts-rule-member-seq-range-minus-1 ()
  rb-compile-simplification-equality
  member-seq-range-minus-1

```

```

function ts-rule-member-plus-1-seq-range ()
  rb-compile-simplification-equality
  member-plus-1-seq-range

```

```

function ts-rule-definedness-of-empty-inverse-map ()
  rb-compile-simplification-equality
  definedness-of-empty-inverse-map

```

```

function ts-rule-universal-lower-bound-of-less-constrained-than ()
  rb-compile-simplification-equality
  universal-lower-bound-of-less-constrained-than

```

```

function TS-RULE-EXTRACT-A-SCHEDULE-PRODUCES-AN-INSTANCE-OF-INVERSE-MAP ()
  rb-compile-simplification-equality
  EXTRACT-A-SCHEDULE-PRODUCES-AN-INSTANCE-OF-INVERSE-MAP

```

```

function TS-RULE-INVERSE-MAP-SATISFIES-IMPLIES-P-CONSISTENT-TIME-INTERVALS ()
  rb-compile-implication-forward

```

INVERSE-MAP-SATISFIES-IMPLIES-P-CONSISTENT-TIME-INTERVALS

```
function TS-RULE-CONSISTENT-POE-to-P-CONSISTENT-POE ()  
  rb-compile-implication-forward  
  CONSISTENT-POE-to-P-CONSISTENT-POE
```

```
function TS-RULE-CONSISTENT-POD-to-P-CONSISTENT-POD ()  
  rb-compile-implication-forward  
  CONSISTENT-POD-to-P-CONSISTENT-POD
```

```
function  
  TS-RULE-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-to-P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE  
  ()  rb-compile-implication-forward  
  CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-to-P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE
```

```
function TS-RULE-CONSISTENT-PAX-RESOURCE-CAPACITY-to-P-CONSISTENT-PAX-RESOURCE-CAPACITY  
  rb-compile-implication-forward  
  CONSISTENT-PAX-RESOURCE-CAPACITY-to-P-CONSISTENT-PAX-RESOURCE-CAPACITY
```

```
function  
  TS-RULE-CONSISTENT-CARGO-RESOURCE-CAPACITY-to-P-CONSISTENT-CARGO-RESOURCE-CAPACITY ()  
  rb-compile-implication-forward  
  CONSISTENT-CARGO-RESOURCE-CAPACITY-to-P-CONSISTENT-CARGO-RESOURCE-CAPACITY
```

```
function TS-RULE-AVAILABLE-RESOURCES-USED-to-P-AVAILABLE-RESOURCES-USED ()  
  rb-compile-implication-forward  
  AVAILABLE-RESOURCES-USED-to-P-AVAILABLE-RESOURCES-USED
```

```
function TS-RULE-SCHEDULED-MVRS-to-P-SCHEDULED-MVRS ()  
  rb-compile-implication-forward  
  SCHEDULED-MVRS-to-P-SCHEDULED-MVRS
```

```
function TS-RULE-CONSISTENT-POE-over-extract ()  
  rb-compile-simplification-equality  
  CONSISTENT-POE-over-extract
```

```
function TS-RULE-CONSISTENT-POD-over-extract ()  
  rb-compile-simplification-equality  
  CONSISTENT-POD-over-extract
```

```
function TS-RULE-CONSISTENT-RELEASE-TIMES-over-extract ()
```

```

rb-compile-implication-backward
CONSISTENT-RELEASE-TIMES-over-extract

function TS-RULE-CONSISTENT-DUE-TIMES-over-extract ()
  rb-compile-implication-backward
  CONSISTENT-DUE-TIMES-over-extract

function TS-RULE-CONSISTENT-TRIP-SEPARATION-over-extract ()
  rb-compile-implication-backward
  CONSISTENT-TRIP-SEPARATION-over-extract

function TS-RULE-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-over-extract ()
  rb-compile-simplification-equality
  CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-over-extract

function TS-RULE-CONSISTENT-PAX-RESOURCE-CAPACITY-over-extract ()
  rb-compile-simplification-equality
  CONSISTENT-PAX-RESOURCE-CAPACITY-over-extract

function
  TS-RULE-CONSISTENT-CARGO-RESOURCE-CAPACITY-over-extract ()
  rb-compile-simplification-equality
  CONSISTENT-CARGO-RESOURCE-CAPACITY-over-extract

function TS-RULE-AVAILABLE-RESOURCES-USED-over-extract ()
  rb-compile-simplification-equality
  AVAILABLE-RESOURCES-USED-over-extract

function TS-RULE-SCHEDULED-MVRS-over-extract ()
  rb-compile-simplification-equality
  SCHEDULED-MVRS-over-extract

function TS-RULE-DISTRIBUTE-P-CONSISTENT-POE-OVER-EMPTY-INVERSE-MAP ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-POE-OVER-EMPTY-INVERSE-MAP

function TS-RULE-DISTRIBUTE-P-CONSISTENT-POE-OVER-INVERSE-MAP-WITH! ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-POE-OVER-INVERSE-MAP-WITH!

function TS-RULE-DISTRIBUTE-P-CONSISTENT-POE-OVER-INVERSE-MAP-APPEND ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-POE-OVER-INVERSE-MAP-APPEND

```

```
function TS-RULE-DISTRIBUTE-P-CONSISTENT-POE-OVER-INVERSE-MAP-UNION! ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-POE-OVER-INVERSE-MAP-UNION!
```

```
function TS-RULE-DISTRIBUTE-P-CONSISTENT-POD-OVER-EMPTY-INVERSE-MAP ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-POD-OVER-EMPTY-INVERSE-MAP
```

```
function TS-RULE-DISTRIBUTE-P-CONSISTENT-POD-OVER-INVERSE-MAP-WITH! ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-POD-OVER-INVERSE-MAP-WITH!
```

```
function TS-RULE-DISTRIBUTE-P-CONSISTENT-POD-OVER-INVERSE-MAP-APPEND ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-POD-OVER-INVERSE-MAP-APPEND
```

```
function TS-RULE-DISTRIBUTE-P-CONSISTENT-POD-OVER-INVERSE-MAP-UNION! ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-POD-OVER-INVERSE-MAP-UNION!
```

```
function TS-RULE-DISTRIBUTE-P-CONSISTENT-PAX-RESOURCE-CAPACITY-OVER-EMPTY-INVERSE-MAP ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-PAX-RESOURCE-CAPACITY-OVER-EMPTY-INVERSE-MAP
```

```
function
  TS-RULE-DISTRIBUTE-P-CONSISTENT-PAX-RESOURCE-CAPACITY-OVER-INVERSE-MAP-WITH! ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-PAX-RESOURCE-CAPACITY-OVER-INVERSE-MAP-WITH!
```

```
function
  TS-RULE-DISTRIBUTE-P-CONSISTENT-PAX-RESOURCE-CAPACITY-OVER-INVERSE-MAP-APPEND ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-PAX-RESOURCE-CAPACITY-OVER-INVERSE-MAP-APPEND
```

```
function TS-RULE-DISTRIBUTE-P-CONSISTENT-PAX-RESOURCE-CAPACITY-OVER-INVERSE-MAP-UNION! ()
  rb-compile-implication-forward
  DISTRIBUTE-P-CONSISTENT-PAX-RESOURCE-CAPACITY-OVER-INVERSE-MAP-UNION!
```

```
function TS-RULE-DISTRIBUTE-P-CONSISTENT-CARGO-RESOURCE-CAPACITY-OVER-EMPTY-INVERSE-MAP
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-CARGO-RESOURCE-CAPACITY-OVER-EMPTY-INVERSE-MAP
```

```

function
  TS-RULE-DISTRIBUTE-P-CONSISTENT-CARGO-RESOURCE-CAPACITY-OVER-INVERSE-MAP-WITH! ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-CARGO-RESOURCE-CAPACITY-OVER-INVERSE-MAP-WITH!

function
  TS-RULE-DISTRIBUTE-P-CONSISTENT-CARGO-RESOURCE-CAPACITY-OVER-INVERSE-MAP-APPEND ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-CARGO-RESOURCE-CAPACITY-OVER-INVERSE-MAP-APPEND

function TS-RULE-DISTRIBUTE-P-CONSISTENT-CARGO-RESOURCE-CAPACITY-OVER-INVERSE-MAP-UNION!
  rb-compile-implication-forward
  DISTRIBUTE-P-CONSISTENT-CARGO-RESOURCE-CAPACITY-OVER-INVERSE-MAP-UNION!

function
  TS-RULE-DISTRIBUTE-P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-OVER-EMPTY-INVERSE-MAP ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-OVER-EMPTY-INVERSE-MAP

function
  TS-RULE-DISTRIBUTE-P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-OVER-INVERSE-MAP-WITH! ()
  rb-compile-simplification-equality
  DISTRIBUTE-p-consistent-movement-type-and-resource-OVER-INVERSE-MAP-WITH!

function
  TS-RULE-DISTRIBUTE-P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-OVER-INVERSE-MAP-APPEND ()
  rb-compile-simplification-equality
  DISTRIBUTE-p-consistent-movement-type-and-resource-OVER-INVERSE-MAP-APPEND

function TS-RULE-DISTRIBUTE-P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE-OVER-INVERSE-MAP-UNION!
  rb-compile-simplification-equality
  DISTRIBUTE-p-consistent-movement-type-and-resource-OVER-INVERSE-MAP-UNION!

function TS-RULE-DISTRIBUTE-P-AVAILABLE-RESOURCES-USED-OVER-EMPTY-INVERSE-MAP ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-AVAILABLE-RESOURCES-USED-OVER-EMPTY-INVERSE-MAP

function TS-RULE-DISTRIBUTE-P-AVAILABLE-RESOURCES-USED-OVER-INVERSE-MAP-WITH! ()
  rb-compile-simplification-equality
  DISTRIBUTE-P-AVAILABLE-RESOURCES-USED-OVER-INVERSE-MAP-WITH!

function TS-RULE-DISTRIBUTE-P-AVAILABLE-RESOURCES-USED-OVER-INVERSE-MAP-APPEND ()

```

```

rb-compile-simplification-equality
DISTRIBUTE-P-AVAILABLE-RESOURCES-USED-OVER-INVERSE-MAP-APPEND

function TS-RULE-DISTRIBUTE-P-AVAILABLE-RESOURCES-USED-OVER-INVERSE-MAP-UNION! ()
rb-compile-simplification-equality
DISTRIBUTE-P-AVAILABLE-RESOURCES-USED-OVER-INVERSE-MAP-UNION!

function TS-RULE-DISTRIBUTE-ALL-MVRS-P-SCHEDULED-OVER-EMPTY-INVERSE-MAP ()
rb-compile-simplification-equality
DISTRIBUTE-ALL-MVRS-P-SCHEDULED-OVER-EMPTY-INVERSE-MAP

function TS-RULE-PAX-RESOURCE-USED-IN-MVRS-OVER-SINGLETON-MANIFEST ()
rb-compile-simplification-equality
PAX-RESOURCE-USED-IN-MVRS-OVER-SINGLETON-MANIFEST

function TS-RULE-CARGO-RESOURCE-USED-IN-MVRS-OVER-SINGLETON-MANIFEST ()
rb-compile-simplification-equality
CARGO-RESOURCE-USED-IN-MVRS-OVER-SINGLETON-MANIFEST

function TS-RULE-DISTRIBUTE-P-CONSISTENT-TIME-INTERVALS-OVER-EMPTY-INVERSE-MAP ()
rb-compile-simplification-equality
DISTRIBUTE-P-CONSISTENT-TIME-INTERVALS-OVER-EMPTY-INVERSE-MAP

function TS-RULE-DISTRIBUTE-P-CONSISTENT-TIME-INTERVALS-OVER-INVERSE-MAP-WITH! ()
rb-compile-simplification-equality
DISTRIBUTE-P-CONSISTENT-TIME-INTERVALS-OVER-INVERSE-MAP-WITH!

function TS-RULE-DISTRIBUTE-P-CONSISTENT-TIME-INTERVALS-OVER-INVERSE-MAP-APPEND ()
rb-compile-simplification-equality
DISTRIBUTE-P-CONSISTENT-TIME-INTERVALS-OVER-INVERSE-MAP-APPEND

function TS-RULE-SELECT-MANIFEST-FROM-LITERAL-TRIP ()
rb-compile-simplification-equality
SELECT-MANIFEST-FROM-LITERAL-TRIP

function TS-RULE-SELECT-MANIFEST-FROM-LITERAL-P-TRIP ()
rb-compile-simplification-equality
SELECT-MANIFEST-FROM-LITERAL-P-TRIP

```

THEORY-MISC-DEFS

```
function TS-RULE-CONSISTENT-RELEASE-TIMES-TO-P-CONSISTENT-RELEASE-TIMES (a)
  computed-using
  a = 'fa(S0)(inverse-map-satisfies(S0,@PS0) => CONSISTENT-RELEASE-TIMES(S0))'
  & new-a = make-structure(
    '##r RB-GRAMMAR
    (rule-instance-make UNDEFINED,
     P-CONSISTENT-RELEASE-TIMES(@c-t(PS0))),
    $ {}, $ {},
    1, TS-RULE-CONSISTENT-RELEASE-TIMES-TO-P-CONSISTENT-RELEASE-TIMES)'
  => TS-RULE-CONSISTENT-RELEASE-TIMES-TO-P-CONSISTENT-RELEASE-TIMES(a) = new-a

function TS-RULE-CONSISTENT-DUE-TIMES-TO-P-CONSISTENT-DUE-TIMES (a)
  computed-using
  a = 'fa(S0)(inverse-map-satisfies(S0,@PS0) => CONSISTENT-DUE-TIMES(S0))'
  & new-a = make-structure(
    '##r RB-GRAMMAR
    (rule-instance-make UNDEFINED,
     P-CONSISTENT-DUE-TIMES(@c-t(PS0))),
    $ {}, $ {},
    1, TS-RULE-CONSISTENT-DUE-TIMES-TO-P-CONSISTENT-DUE-TIMES)'
  => TS-RULE-CONSISTENT-DUE-TIMES-TO-P-CONSISTENT-DUE-TIMES(a) = new-a

function TS-RULE-CONSISTENT-TRIP-SEPARATION-TO-P-CONSISTENT-TRIP-SEPARATION (a)
  computed-using
  a = 'fa(S0)(inverse-map-satisfies(S0,@PS0) => CONSISTENT-TRIP-SEPARATION(S0))'
  & new-a = make-structure(
    '##r RB-GRAMMAR
    (rule-instance-make UNDEFINED,
     P-CONSISTENT-TRIP-SEPARATION-EST(@c-t(PS0))
     & P-CONSISTENT-TRIP-SEPARATION-LST(@c-t(PS0))),
    $ {}, $ {},
    1, TS-RULE-CONSISTENT-TRIP-SEPARATION-TO-P-CONSISTENT-TRIP-SEPARATION)'
  => TS-RULE-CONSISTENT-TRIP-SEPARATION-TO-P-CONSISTENT-TRIP-SEPARATION(a) = new-a

end-theory
```

C Global Search Theory for Transportation Scheduling

(Global-Search-Theory GS-CP-TRANSPORTATION-SCHEDULING

```
input-types seq(movement-record),
             seq(resource-name)
input-vars gs-mvrs,
           gs-assets
input-condition true

output-types schedule
output-vars sched
output-condition
  Available-Resources-Used(gs-assets, sched)
  & Scheduled-Mvrs(sched) = seq-to-set(gs-mvrs)

subspace-types
  p-schedule,
  seq(movement-record),
  seq(movement-record)
subspace-vars
  p-sched,
  sched-mvrs,
  unsched-mvrs
subspace-split-vars
  p-sched-new,
  sched-mvrs-new,
  unsched-mvrs-new
subspace-vars-constraint
  inverse-map?(gs-assets, p-sched)
  & seq-to-set(sched-mvrs) = p-scheduled-mvrs(p-sched)
  & seq-to-set(gs-mvrs)
    = seq-to-set(sched-mvrs) union seq-to-set(unsched-mvrs)
  & disjoint(seq-to-set(sched-mvrs), seq-to-set(unsched-mvrs))
  & defined?(p-sched)

Constraint-Info-types
  resource-name, integer, movement-record

Constraint-Info-vars
  gs-resource, gs-trp, gs-mvr

Constraint-Info-condition
```

```

(((gs-resource in gs-assets
  & gs-trp in [1 .. size(p-sched(gs-resource))])
 or* (gs-resource in gs-assets
  & gs-trp = size(p-sched(gs-resource)) + 1))
 & gs-mvr = first(unsched-mvrs))

Splitting-constraint
sched-mvrs-new = prepend(sched-mvrs, gs-mvr)
& unsched-mvrs-new = rest(unsched-mvrs)
&
(if gs-trp in [1 .. size(p-sched(gs-resource))])
then p-sched-new(gs-resource)(gs-trp).manifest
  = prepend(p-sched(gs-resource)(gs-trp).manifest, gs-mvr)
else
  p-sched-new(gs-resource)
  = append(p-sched(gs-resource),
    <gs-mvr.poe, gs-mvr.pod, gs-mvr.release-date,
    gs-mvr.due-date
    - duration(gs-mvr.distance,
    *resource-map*(gs-resource)),
    duration(gs-mvr.distance,
    *resource-map*(gs-resource)),
    [gs-mvr]>) )

satisfies
  inverse-map-satisfies(sched, p-sched)
refines
  less-constrained-than(p-sched, p-sched-new)

initial-space
  (<empty-inverse-map(gs-assets), [], gs-mvrs>)

extract
  sched = extract-a-schedule(p-sched)
)

```

D Derived Transportation Scheduler

Derivation Summary

0. Focus Initialize TS
1. Tactic Global Search with Constraint Propagation on TS
2. Remove foci TS-PROPAGATE, TS-CUT
3. Simplify, context-independent-fast:
 if ## then ## else some (SCHED-2: #...
4. Simplify, context-dependent, forward-0, backward-4: if ## else und...
5. Simplify, context-dependent, forward-1, backward-5: ## & ## & ## &...
6. Simplify, context-dependent, forward-0, backward-4: seq-to-set(##)...
7. Simplify, context-independent-fast: true & UNSCHED-MVRS = []
8. Abstract into LET:
 let (var Z-109 = some (SCHED-2: map(RESOURCE-NAME, seq(T...
9. Simplify, context-dependent, forward-0, backward-4: ## & ## & ## &...
10. Abstract TS-PROPAGATE-WITH!(P-SCHED, R, P-TRP, MVR)
 into PROPAGATED-P-SCHED in ex (P-TRP: ##, R: ##) (## &...
11. Simplify, context-dependent, forward-0, backward-4: ## & ## & ## &...
12. Abstract TS-PROPAGATE-APPEND(P-SCHED, R, <##, ##, #...
 into PROPAGATED-P-SCHED in ex (R: ##) (## & ## & ## & ...
13. Simplify, context-independent-fast: if ## & ## ...
14. Refine compile into Lisp: TS-AUX
15. Refine compile into Lisp: TS

Initial Specification

```
function TS
  (mvrs : seq(movement-record),
   assets : seq(resource-name) | true)
  returns ( sched : map( resource-name, seq(trip) )

    | Consistent-POE(sched)
    & Consistent-POD(sched)
    & Consistent-Release-Times(sched)
    & Consistent-Arrival-Times(sched)
    & Consistent-Due-Times(sched)
    & Consistent-Trip-Separation(sched)
    & Consistent-Pax-Resource-Capacity(sched)
    & Consistent-Cargo-Resource-Capacity(sched)
    & Consistent-Movement-Type-and-Resource(sched)
    & Available-Resources-Used(assets, sched)
    & Scheduled-mvrs(sched) = seq-to-set(mvrs)
  )
```

Derived Code

The following is the derived Refine code of the scheduler. The first two functions, TS-AUX and TS, are the backtrack scheduler and these are machine generated in KIDS. The remaining functions perform constraint propagation when a movement-record is added to an existing trip or a newly created trip respectively.

TS-SPLIT-AND-PROPAGATE encasulates the subspace splitting operation and resulting propagation (and pruning). TS-SPLIT-AND-PROPAGATE-1 adds a movement-record to an existing trip and propagates the effect. TS-SPLIT-AND-PROPAGATE-2 adds a movement-record to a newly created trip and propagates the effect. PROPAGATE-LATEST-START-TIME propagates the effect of a change to the latest start time of a trip. Similarly, PROPAGATE-EARLIEST-START-TIME propagates the effect of a change to the earliest start time of a trip.

```
function TS
  (MVRS: seq(MOVEMENT-RECORD), ASSETS: seq(RESOURCE-NAME))
  returns
    (SCHED: map(RESOURCE-NAME, seq(TRIP))
     | CONSISTENT-POE(SCHED) & CONSISTENT-POD(SCHED)
     & CONSISTENT-RELEASE-TIMES(SCHED)
     & CONSISTENT-DUE-TIMES(SCHED)
     & CONSISTENT-TRIP-SEPARATION(SCHED)
     & CONSISTENT-PAX-RESOURCE-CAPACITY(SCHED)
     & CONSISTENT-CARGO-RESOURCE-CAPACITY(SCHED)
     & CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE(SCHED)
     & AVAILABLE-RESOURCES-USED(ASSETS, SCHED)
     & SCHEDULED-MVRS(SCHED) = seq-to-set(MVRS))
  = let (NEW-ROHAT-1
        : tuple
          (P-SCHEDULE, seq(MOVEMENT-RECORD), seq(MOVEMENT-RECORD))
          = TS-INITIAL-PROPAGATE(MVRS, ASSETS))
        if DEFINED?(NEW-ROHAT-1)
        then TS-AUX
          (MVRS, ASSETS, NEW-ROHAT-1.1, NEW-ROHAT-1.2, NEW-ROHAT-1.3)
```

```
function TS-AUX
  (MVRS: seq(MOVEMENT-RECORD), ASSETS: seq(RESOURCE-NAME),
   P-SCHED: P-SCHEDULE, SCHED-MVRS: seq(MOVEMENT-RECORD),
   UNSCHED-MVRS: seq(MOVEMENT-RECORD))
  | INVERSE-MAP?(ASSETS, P-SCHED)
  & seq-to-set(SCHED-MVRS) = P-SCHEDULED-MVRS(P-SCHED)
  & seq-to-set(MVRS)
    = seq-to-set(UNSCHED-MVRS) union seq-to-set(SCHED-MVRS)
  & DISJOINT(seq-to-set(SCHED-MVRS), seq-to-set(UNSCHED-MVRS))
```

```

& DEFINED?(P-SCHED) & P-CONSISTENT-RELEASE-TIMES(P-SCHED)
& P-CONSISTENT-DUE-TIMES(P-SCHED)
& P-CONSISTENT-TRIP-SEPARATION-EST(P-SCHED)
& P-CONSISTENT-TRIP-SEPARATION-LST(P-SCHED)
& P-CONSISTENT-POE(P-SCHED) & P-CONSISTENT-POD(P-SCHED)
& P-CONSISTENT-PAX-RESOURCE-CAPACITY(P-SCHED)
& P-CONSISTENT-CARGO-RESOURCE-CAPACITY(P-SCHED)
& P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE(P-SCHED)
& P-AVAILABLE-RESOURCES-USED(ASSETS, P-SCHED))
: map(RESOURCE-NAME, seq(TRIP))
= if UNSCHED-MVRS = [] then EXTRACT-A-SCHEDULE(P-SCHED)
  else OTHERWISE
    (some (SCHED-2: map(RESOURCE-NAME, seq(TRIP)))
      ex (C-26:
          tuple
            (P-SCHEDULE, seq(MOVEMENT-RECORD), seq(MOVEMENT-RECORD)),
            GS-RESOURCE: RESOURCE-NAME, GS-TRP: integer)
          (GS-TRP in [1 .. size(P-SCHED(GS-RESOURCE))]
            & GS-RESOURCE in ASSETS & DEFINED?(SCHED-2)
            & SCHED-2
              = TS-AUX
                (MVRS, ASSETS, C-26.1, C-26.2, C-26.3)
            & DEFINED?(C-26)
            & C-26
              = TS-SPLIT-AND-PROPAGATE
                (MVRS, ASSETS, P-SCHED, SCHED-MVRS, UNSCHED-MVRS,
                  GS-RESOURCE, GS-TRP, first(UNSCHED-MVRS))),
          some (SCHED-2: map(RESOURCE-NAME, seq(TRIP)))
            ex (C-27:
                tuple
                  (P-SCHEDULE, seq(MOVEMENT-RECORD), seq(MOVEMENT-RECORD)),
                  GS-RESOURCE: RESOURCE-NAME)
                (GS-RESOURCE in ASSETS & DEFINED?(SCHED-2)
                  & SCHED-2
                    = TS-AUX
                      (MVRS, ASSETS, C-27.1, C-27.2, C-27.3)
                  & DEFINED?(C-27)
                  & C-27
                    = TS-SPLIT-AND-PROPAGATE
                      (MVRS, ASSETS, P-SCHED, SCHED-MVRS, UNSCHED-MVRS,
                        GS-RESOURCE, 1 + size(P-SCHED(GS-RESOURCE)),
                        first(UNSCHED-MVRS))))))

```

```

function TS-INITIAL-PROPAGATE
  (MVRS: seq(MOVEMENT-RECORD), ASSETS: seq(RESOURCE-NAME))

```

```

returns
(S-HAT-2:
  tuple
    (P-SCHEDULE, seq(MOVEMENT-RECORD), seq(MOVEMENT-RECORD))
  | EXTREMAL
    (S-HAT-2,
     lambda (var T1, var T2) LESS-CONSTRAINED-THAN(T1.1, T2.1),
     {T-HAT-2
      | (T-HAT-2:
         tuple
           (P-SCHEDULE, seq(MOVEMENT-RECORD), seq(MOVEMENT-RECORD)))
         LESS-CONSTRAINED-THAN(EMPTY-INVERSE-MAP(ASSETS), T-HAT-2.1)
         & P-CONSISTENT-RELEASE-TIMES(T-HAT-2.1)
         & P-CONSISTENT-DUE-TIMES(T-HAT-2.1)
         & P-CONSISTENT-TRIP-SEPARATION-EST(T-HAT-2.1)
         & P-CONSISTENT-TRIP-SEPARATION-LST(T-HAT-2.1)
         & P-CONSISTENT-POE(T-HAT-2.1) & P-CONSISTENT-POD(T-HAT-2.1)
         & P-CONSISTENT-PAX-RESOURCE-CAPACITY(T-HAT-2.1)
         & P-CONSISTENT-CARGO-RESOURCE-CAPACITY(T-HAT-2.1)
         & P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE(T-HAT-2.1)
         & P-AVAILABLE-RESOURCES-USED(ASSETS, T-HAT-2.1)}))
    = <empty-inverse-map(assets), [], MVRs>

```

```

function TS-SPLIT-AND-PROPAGATE
(MVRS: seq(MOVEMENT-RECORD), ASSETS: seq(RESOURCE-NAME),
 P-SCHED: P-SCHEDULE, SCHED-MVRS: seq(MOVEMENT-RECORD),
 UNSCHED-MVRS: seq(MOVEMENT-RECORD),
 GS-RESOURCE: RESOURCE-NAME, GS-TRP: integer,
 GS-MVR: MOVEMENT-RECORD
 | INVERSE-MAP?(ASSETS, P-SCHED)
   & seq-to-set(SCHED-MVRS) = P-SCHEDULED-MVRS(P-SCHED)
   & seq-to-set(MVRS)
     = seq-to-set(UNSCHED-MVRS) union seq-to-set(SCHED-MVRS)
   & DISJOINT(seq-to-set(SCHED-MVRS), seq-to-set(UNSCHED-MVRS))
   & DEFINED?(P-SCHED)
   & (GS-TRP in [1 .. size(P-SCHED(GS-RESOURCE))])
     & GS-RESOURCE in ASSETS
     or* GS-TRP = 1 + size(P-SCHED(GS-RESOURCE))
       & GS-RESOURCE in ASSETS)
   & GS-MVR = first(UNSCHED-MVRS)
   & P-CONSISTENT-RELEASE-TIMES(P-SCHED)
   & P-CONSISTENT-DUE-TIMES(P-SCHED)
   & P-CONSISTENT-TRIP-SEPARATION-EST(P-SCHED)
   & P-CONSISTENT-TRIP-SEPARATION-LST(P-SCHED)
   & P-CONSISTENT-POE(P-SCHED) & P-CONSISTENT-POD(P-SCHED)
   & P-CONSISTENT-PAX-RESOURCE-CAPACITY(P-SCHED)

```

```

    & P-CONSISTENT-CARGO-RESOURCE-CAPACITY(P-SCHED)
    & P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE(P-SCHED)
    & P-AVAILABLE-RESOURCES-USED(ASSETS, P-SCHED))
returns
(S-HAT-1:
  tuple
    (P-SCHEDULE, seq(MOVEMENT-RECORD), seq(MOVEMENT-RECORD))
  | EXTREMAL
    (S-HAT-1,
     lambda (var T1, var T2) LESS-CONSTRAINED-THAN(T1.1, T2.1),
     {T-HAT-1
      | (T-HAT-1:
         tuple
           (P-SCHEDULE, seq(MOVEMENT-RECORD), seq(MOVEMENT-RECORD)))
        LESS-CONSTRAINED-THAN(P-SCHED, T-HAT-1.1)
        & T-HAT-1.2 = prepend(SCHED-MVRS, GS-MVR)
        & T-HAT-1.3 = rest(UNSCHED-MVRS)
        & (if GS-TRP in [1 .. size(P-SCHED(GS-RESOURCE))])
           then T-HAT-1.1(GS-RESOURCE)(GS-TRP).MANIFEST
              = prepend(P-SCHED(GS-RESOURCE)(GS-TRP).MANIFEST, GS-MVR)
           else T-HAT-1.1(GS-RESOURCE)
              = append
                (P-SCHED(GS-RESOURCE),
                 <GS-MVR.POE, GS-MVR.POD, GS-MVR.RELEASE-DATE,
                 GS-MVR.DUE-DATE
                  - DURATION(GS-MVR.DISTANCE, *RESOURCE-MAP*(GS-RESOURCE)
                  DURATION(GS-MVR.DISTANCE, *RESOURCE-MAP*(GS-RESOURCE))
                 [GS-MVR]>))
        & P-CONSISTENT-RELEASE-TIMES(T-HAT-1.1)
        & P-CONSISTENT-DUE-TIMES(T-HAT-1.1)
        & P-CONSISTENT-TRIP-SEPARATION-EST(T-HAT-1.1)
        & P-CONSISTENT-TRIP-SEPARATION-LST(T-HAT-1.1)
        & P-CONSISTENT-POE(T-HAT-1.1) & P-CONSISTENT-POD(T-HAT-1.1)
        & P-CONSISTENT-PAX-RESOURCE-CAPACITY(T-HAT-1.1)
        & P-CONSISTENT-CARGO-RESOURCE-CAPACITY(T-HAT-1.1)
        & P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE(T-HAT-1.1)
        & P-AVAILABLE-RESOURCES-USED(ASSETS, T-HAT-1.1)}))
= if GS-TRP in [1 .. size(P-SCHED(GS-RESOURCE))])
  then TS-SPLIT-AND-PROPAGATE-1
    (MVRS, ASSETS, P-SCHED, SCHED-MVRS, UNSCHED-MVRS,
     GS-RESOURCE, GS-TRP, GS-MVR)
  else TS-SPLIT-AND-PROPAGATE-2
    (MVRS, ASSETS, P-SCHED, SCHED-MVRS, UNSCHED-MVRS,
     GS-RESOURCE, GS-TRP, GS-MVR)

```

```
function PROPAGATE-LATEST-START-TIME-1
```

```

(CHANGE-TUPLE-3:
  tuple(RESOURCE-NAME, integer, integer),
  STATE: tuple(P-SCHEDULE, seq(RESOURCE-NAME))
  | ex (ASSETS: seq(RESOURCE-NAME), P-SCHED-N: P-SCHEDULE,
        Y-1: integer, P-TRP: integer, RSRC: RESOURCE-NAME)
        (CHANGE-TUPLE-3 = <RSRC, P-TRP, Y-1>
          & STATE = <P-SCHED-N, ASSETS>))
: tuple(P-SCHEDULE, seq(RESOURCE-NAME))
= letp (ASSETS: seq(RESOURCE-NAME), P-SCHED-N: P-SCHEDULE,
        Y-1-P: integer, P-TRP-P: integer,
        RSRC-P: RESOURCE-NAME
        | CHANGE-TUPLE-3 = <RSRC-P, P-TRP-P, Y-1-P>
          & STATE = <P-SCHED-N, ASSETS>)
let (C-3: seq(P-TRIP) = P-SCHED-N(RSRC-P))
let (C-1: P-TRIP = C-3(P-TRP-P))
if C-1.LATEST-START-TIME <= Y-1-P
  then <P-SCHED-N, ASSETS>
  else
  if Y-1-P < C-1.EARLIEST-START-TIME
    then undefined
    else
    let (STATE: tuple(P-SCHEDULE, seq(RESOURCE-NAME))
        = <(P-SCHED-N
            ** {| RSRC-P
                -> SEQ-SHADOW1
                    (C-3, P-TRP-P,
                     TUPLE-SHADOW(C-1.LATEST-START-TIME, Y-1-P))
                |}),
        ASSETS>)
    let (NEW-STATE-SEQ
        : seq(tuple(P-SCHEDULE, seq(RESOURCE-NAME)))
        = [PROPAGATE-LATEST-START-TIME-1
            (<RSRC-P, C-4, C-0>, STATE)
            | (C-4: integer, C-2: P-TRIP,
              C-0: integer)
              C-4 in [1 .. -1 + size(C-3)]
              & C-0 < C-2.LATEST-START-TIME
              & C-0 = Y-1-P - 2 * C-2.P-TRIP-DURATION
              & C-2 = C-3(C-4)
              & C-4 = -1 + P-TRP-P])
        if NEW-STATE-SEQ = [] then STATE else first(NEW-STATE-SEQ)

```

```
function PROPAGATE-EARLIEST-START-TIME-1
```

```

(CHANGE-TUPLE-4:
  tuple(RESOURCE-NAME, integer, integer),
  STATE: tuple(P-SCHEDULE, seq(RESOURCE-NAME))

```

```

| ex (ASSETS: seq(RESOURCE-NAME), P-SCHED-N: P-SCHEDULE,
      Y-0: TIME, P-TRP: integer, RSRC: RESOURCE-NAME)
  (CHANGE-TUPLE-4 = <RSRC, P-TRP, Y-0>
    & STATE = <P-SCHED-N, ASSETS>))
: tuple(P-SCHEDULE, seq(RESOURCE-NAME))
= letp (ASSETS: seq(RESOURCE-NAME), P-SCHED-N: P-SCHEDULE,
        Y-0-P: TIME, P-TRP-P: integer, RSRC-P: RESOURCE-NAME
        | CHANGE-TUPLE-4 = <RSRC-P, P-TRP-P, Y-0-P>
          & STATE = <P-SCHED-N, ASSETS>)
let (C-7: seq(P-TRIP) = P-SCHED-N(RSRC-P))
let (C-6: P-TRIP = C-7(P-TRP-P))
if Y-0-P <= C-6.EARLIEST-START-TIME
  then <P-SCHED-N, ASSETS>
  else
  if C-6.LATEST-START-TIME < Y-0-P
    then undefined
    else
    let (STATE: tuple(P-SCHEDULE, seq(RESOURCE-NAME))
        = <(P-SCHED-N
            +* { | RSRC-P
                -> SEQ-SHADOW1
                  (C-7, P-TRP-P,
                    TUPLE-SHADOW(C-6.EARLIEST-START-TIME, Y-0-P))
              | }),
          ASSETS>)
    let (NEW-STATE-SEQ
        : seq(tuple(P-SCHEDULE, seq(RESOURCE-NAME)))
        = [PROPAGATE-EARLIEST-START-TIME-1
          (<RSRC-P, C-8, C-5>, STATE)
          | (C-8: integer, C-5: integer)
          P-TRP-P in [1 .. -1 + size(C-7)]
          & C-7(C-8).EARLIEST-START-TIME < C-5
          & C-5 = Y-0-P + 2 * C-6.P-TRIP-DURATION
          & C-8 = P-TRP-P + 1])
    if NEW-STATE-SEQ = [] then STATE else first(NEW-STATE-SEQ)

```

function TS-SPLIT-AND-PROPAGATE-1

```

(MVRS: seq(MOVEMENT-RECORD), ASSETS: seq(RESOURCE-NAME),
P-SCHED: P-SCHEDULE, SCHED-MVRS: seq(MOVEMENT-RECORD),
UNSCHED-MVRS: seq(MOVEMENT-RECORD),
GS-RESOURCE: RESOURCE-NAME, GS-TRP: integer,
GS-MVR: MOVEMENT-RECORD
| GS-TRP in [1 .. size(P-SCHED(GS-RESOURCE))]
  & (true or* GS-TRP = 1 + size(P-SCHED(GS-RESOURCE)))
  & INVERSE-MAP?(ASSETS, P-SCHED)
  & seq-to-set(SCHED-MVRS) = P-SCHEDULED-MVRS(P-SCHED)

```

```

& seq-to-set(MVRS)
  = seq-to-set(UNSCHED-MVRS) union seq-to-set(SCHED-MVRS)
& DISJOINT(seq-to-set(SCHED-MVRS), seq-to-set(UNSCHED-MVRS))
& DEFINED?(P-SCHED) & GS-MVR = first(UNSCHED-MVRS)
& P-CONSISTENT-RELEASE-TIMES(P-SCHED)
& P-CONSISTENT-DUE-TIMES(P-SCHED)
& P-CONSISTENT-TRIP-SEPARATION-EST(P-SCHED)
& P-CONSISTENT-TRIP-SEPARATION-LST(P-SCHED)
& P-CONSISTENT-POE(P-SCHED) & P-CONSISTENT-POD(P-SCHED)
& P-CONSISTENT-PAX-RESOURCE-CAPACITY(P-SCHED)
& P-CONSISTENT-CARGO-RESOURCE-CAPACITY(P-SCHED)
& P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE(P-SCHED)
& P-AVAILABLE-RESOURCES-USED(ASSETS, P-SCHED)
& GS-RESOURCE in ASSETS)
returns
(S-HAT-1:
  tuple
    (P-SCHEDULE, seq(MOVEMENT-RECORD), seq(MOVEMENT-RECORD))
  | EXTREMAL
    (S-HAT-1,
      lambda (var T1, var T2) LESS-CONSTRAINED-THAN(T1.1, T2.1),
      {T-HAT-1
        | (T-HAT-1:
          tuple
            (P-SCHEDULE, seq(MOVEMENT-RECORD), seq(MOVEMENT-RECORD)))
          T-HAT-1.1(GS-RESOURCE)(GS-TRP).MANIFEST
            = prepend(P-SCHED(GS-RESOURCE)(GS-TRP).MANIFEST, GS-MVR)
            & P-AVAILABLE-RESOURCES-USED(ASSETS, T-HAT-1.1)
            & P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE(T-HAT-1.1)
            & P-CONSISTENT-CARGO-RESOURCE-CAPACITY(T-HAT-1.1)
            & P-CONSISTENT-PAX-RESOURCE-CAPACITY(T-HAT-1.1)
            & P-CONSISTENT-POD(T-HAT-1.1) & P-CONSISTENT-POE(T-HAT-1.1)
            & P-CONSISTENT-TRIP-SEPARATION-LST(T-HAT-1.1)
            & P-CONSISTENT-TRIP-SEPARATION-EST(T-HAT-1.1)
            & P-CONSISTENT-DUE-TIMES(T-HAT-1.1)
            & P-CONSISTENT-RELEASE-TIMES(T-HAT-1.1)
            & T-HAT-1.3 = rest(UNSCHED-MVRS)
            & T-HAT-1.2 = prepend(SCHED-MVRS, GS-MVR)
            & LESS-CONSTRAINED-THAN(P-SCHED, T-HAT-1.1)}))
= let (C-12: seq(P-TRIP) = P-SCHED(GS-RESOURCE))
  let (C-10: P-TRIP = C-12(GS-TRP))
  let (NEW-STATE: tuple(P-SCHEDULE, seq(RESOURCE-NAME))
    = if GS-TRP in [1 .. size(C-12)]
      & GS-MVR.POE ~= C-10.POE
      then undefined
      else

```

```

if GS-TRP in [1 .. size(C-12)]
    & GS-MVR.POD ~= C-10.POD
    then undefined
    else
let (C-11: RESOURCE = *RESOURCE-MAP*(GS-RESOURCE))
if GS-TRP in [1 .. size(C-12)]
    & C-11.RES-CLASS ~in
        CARGO-MODE-TABLE(GS-MVR.MOVE-TYPE, GS-MVR.MODE)
    then undefined
    else
if GS-TRP in [1 .. size(C-12)]
    & C-11.PAX-CAPACITY
        < PAX-RESOURCE-USED-IN-MVRS
            (prepend(C-10.MANIFEST, GS-MVR))
    then undefined
    else
if GS-TRP in [1 .. size(C-12)]
    & C-11.CARGO-CAPACITY
        < CARGO-RESOURCE-USED-IN-MVRS
            (prepend(C-10.MANIFEST, GS-MVR))
    then undefined
    else
let (STATE: tuple(P-SCHEDULE, seq(RESOURCE-NAME))
    = <(P-SCHED
        +* {| GS-RESOURCE
            -> SEQ-SHADOW1
                (C-12, GS-TRP,
                TUPLE-SHADOW
                    (C-10.MANIFEST, prepend(C-10.MANIFEST, GS
                    |}),
                ASSETS>)
let (STATE: tuple(P-SCHEDULE, seq(RESOURCE-NAME))
    = let (NEW-STATE-SEQ
        : seq(tuple(P-SCHEDULE, seq(RESOURCE-NAME)))
        = [PROPAGATE-EARLIEST-START-TIME-1
            (<GS-RESOURCE, GS-TRP, C-13>, STATE)
            | (C-13: TIME)
            GS-TRP in [1 .. size(C-12)]
            & C-10.EARLIEST-START-TIME < C-13
            & C-13 = GS-MVR.RELEASE-DATE])
        if NEW-STATE-SEQ = [] then STATE else first(NEW-STATE-SEQ))
if ~DEFINED?(STATE) then undefined
    else
let (NEW-STATE-SEQ
    : seq(tuple(P-SCHEDULE, seq(RESOURCE-NAME)))
    = [PROPAGATE-LATEST-START-TIME-1

```

```

        (<GS-RESOURCE, GS-TRP, C-9>, STATE)
        | (C-9: integer)
        GS-TRP in [1 .. size(C-12)]
        & C-9 < C-10.LATEST-START-TIME
        & C-9 = GS-MVR.DUE-DATE - C-10.P-TRIP-DURATION])
    if NEW-STATE-SEQ = [] then STATE else first(NEW-STATE-SEQ))
if DEFINED?(NEW-STATE)
  then <NEW-STATE.1, prepend(SCHED-MVRS, GS-MVR),
      rest(UNSCHED-MVRS)>
  else undefined

function TS-SPLIT-AND-PROPAGATE-2
(MVRS: seq(MOVEMENT-RECORD), ASSETS: seq(RESOURCE-NAME),
 P-SCHED: P-SCHEDULE, SCHED-MVRS: seq(MOVEMENT-RECORD),
 UNSCHED-MVRS: seq(MOVEMENT-RECORD),
 GS-RESOURCE: RESOURCE-NAME, GS-TRP: integer,
 GS-MVR: MOVEMENT-RECORD
 | GS-TRP ~in [1 .. size(P-SCHED(GS-RESOURCE))]
 & INVERSE-MAP?(ASSETS, P-SCHED)
 & seq-to-set(SCHED-MVRS) = P-SCHEDULED-MVRS(P-SCHED)
 & seq-to-set(MVRS)
   = seq-to-set(UNSCHED-MVRS) union seq-to-set(SCHED-MVRS)
 & DISJOINT(seq-to-set(SCHED-MVRS), seq-to-set(UNSCHED-MVRS))
 & DEFINED?(P-SCHED)
 & GS-TRP = 1 + size(P-SCHED(GS-RESOURCE))
 & GS-RESOURCE in ASSETS & GS-MVR = first(UNSCHED-MVRS)
 & P-CONSISTENT-RELEASE-TIMES(P-SCHED)
 & P-CONSISTENT-DUE-TIMES(P-SCHED)
 & P-CONSISTENT-TRIP-SEPARATION-EST(P-SCHED)
 & P-CONSISTENT-TRIP-SEPARATION-LST(P-SCHED)
 & P-CONSISTENT-POE(P-SCHED) & P-CONSISTENT-POD(P-SCHED)
 & P-CONSISTENT-PAX-RESOURCE-CAPACITY(P-SCHED)
 & P-CONSISTENT-CARGO-RESOURCE-CAPACITY(P-SCHED)
 & P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE(P-SCHED)
 & P-AVAILABLE-RESOURCES-USED(ASSETS, P-SCHED))
returns
(S-HAT-1:
 tuple
 (P-SCHEDULE, seq(MOVEMENT-RECORD), seq(MOVEMENT-RECORD))
 | EXTREMAL
 (S-HAT-1,
 lambda (var T1, var T2) LESS-CONSTRAINED-THAN(T1.1, T2.1),
 {T-HAT-1
 | (T-HAT-1:
 tuple
 (P-SCHEDULE, seq(MOVEMENT-RECORD), seq(MOVEMENT-RECORD)))

```

```

T-HAT-1.1(GS-RESOURCE)
  = append
    (P-SCHED(GS-RESOURCE),
     <GS-MVR.POE, GS-MVR.POD, GS-MVR.RELEASE-DATE,
     GS-MVR.DUE-DATE
     - DURATION(GS-MVR.DISTANCE, *RESOURCE-MAP*(GS-RESOURCE)),
     DURATION(GS-MVR.DISTANCE, *RESOURCE-MAP*(GS-RESOURCE)),
     [GS-MVR]>)
    & P-AVAILABLE-RESOURCES-USED(ASSETS, T-HAT-1.1)
    & P-CONSISTENT-MOVEMENT-TYPE-AND-RESOURCE(T-HAT-1.1)
    & P-CONSISTENT-CARGO-RESOURCE-CAPACITY(T-HAT-1.1)
    & P-CONSISTENT-PAX-RESOURCE-CAPACITY(T-HAT-1.1)
    & P-CONSISTENT-POD(T-HAT-1.1) & P-CONSISTENT-POE(T-HAT-1.1)
    & P-CONSISTENT-TRIP-SEPARATION-LST(T-HAT-1.1)
    & P-CONSISTENT-TRIP-SEPARATION-EST(T-HAT-1.1)
    & P-CONSISTENT-DUE-TIMES(T-HAT-1.1)
    & P-CONSISTENT-RELEASE-TIMES(T-HAT-1.1)
    & T-HAT-1.3 = rest(UNSCHED-MVRS)
    & T-HAT-1.2 = prepend(SCHED-MVRS, GS-MVR)
    & LESS-CONSTRAINED-THAN(P-SCHED, T-HAT-1.1)))
= let (C-25: RESOURCE = *RESOURCE-MAP*(GS-RESOURCE),
      C-23: seq(MOVEMENT-RECORD) = [GS-MVR])
let (NEW-STATE: tuple(P-SCHEDULE, seq(RESOURCE-NAME))
  = if C-25.PAX-CAPACITY
    < PAX-RESOURCE-USED-IN-MVRS(C-23)
    then undefined
  else
  if C-25.CARGO-CAPACITY
    < CARGO-RESOURCE-USED-IN-MVRS(C-23)
    then undefined
  else
  if C-25.RES-CLASS ~in
    CARGO-MODE-TABLE(GS-MVR.MOVE-TYPE, GS-MVR.MODE)
    then undefined
  else
  let (C-21: integer
      = DURATION(GS-MVR.DISTANCE, C-25))
  let (C-24: seq(P-TRIP) = P-SCHED(GS-RESOURCE),
      C-22: TIME = GS-MVR.RELEASE-DATE,
      C-18: integer = GS-MVR.DUE-DATE - C-21)
  let (STATE: tuple(P-SCHEDULE, seq(RESOURCE-NAME))
      = <(P-SCHED
          +* {| GS-RESOURCE
              -> append
                (C-24,
                 <GS-MVR.POE, GS-MVR.POD, C-22, C-18,

```

```

                                C-21, C-23>)
                                |}),
                                ASSETS>)
let (STATE: tuple(P-SCHEDULE, seq(RESOURCE-NAME))
    = let (NEW-STATE-SEQ
          : seq(tuple(P-SCHEDULE, seq(RESOURCE-NAME)))
          = [PROPAGATE-EARLIEST-START-TIME-1
              (<GS-RESOURCE, C-20 + 1, C-15>, STATE)
              | (C-20: integer, C-16: P-TRIP,
                 C-15: integer)
              1 <= C-20 & C-22 < C-15
              & C-15
              = C-16.EARLIEST-START-TIME
              + 2 * C-16.P-TRIP-DURATION
              & C-16 = C-24(C-20)
              & C-20 = size(C-24)])
          if NEW-STATE-SEQ = [] then STATE else first(NEW-STATE-SEQ))
if ~DEFINED?(STATE) then undefined
else
let (NEW-STATE-SEQ
    : seq(tuple(P-SCHEDULE, seq(RESOURCE-NAME)))
    = [PROPAGATE-LATEST-START-TIME-1
        (<GS-RESOURCE, C-19, C-14>, STATE)
        | (C-19: integer, C-17: P-TRIP,
           C-14: integer)
        1 <= C-19
        & C-14 < C-17.LATEST-START-TIME
        & C-14
        = C-18 - 2 * C-17.P-TRIP-DURATION
        & C-17 = C-24(C-19)
        & C-19 = size(C-24)])
    if NEW-STATE-SEQ = [] then STATE else first(NEW-STATE-SEQ))
if DEFINED?(NEW-STATE)
then <NEW-STATE.1, prepend(SCHED-MVRS, GS-MVR),
    rest(UNSCHED-MVRS)>
else undefined

```