

# A Representational Paradigm for Dynamic Resource Transformation Problems

Warren B. Powell and Joel A. Shapiro

Department of Civil Engineering and Operations Research,  
Princeton University, Princeton, NJ 08544  
Statistics and Operations Research  
Technical Report CL-99-01

March, 1998  
Revised, September, 1998  
Revised, June, 1999  
Revised, August, 1999

DRAFT

The modeling and optimization of complex operational problems represents an emerging opportunity for operations research. Not surprisingly, some of the largest and most complex instances of these problems arise in the management of logistics systems, and especially in the design and control of freight transportation operations. These problems pose a serious challenge in the formulation of mathematical models that capture not only the dynamics of these operations, but also the complex laws that make up the physics of these problems.

We feel that our ability to solve these problems is limited by the languages that we use to express them. Classical mathematical paradigms do not provide an easy and natural way to represent the optimization of these problems under uncertainty, or to capture the complexities of large scale operations. In particular, models do not capture the flow of information in large organizations, preferring instead to assume the presence of a single, all-knowing decision-maker. As a result, most dynamic models posed in the literature are myopic or deterministic.

The characteristics of more complex operations has spawned an extensive literature presenting models that are unique to a particular industry. For example, we solve airline fleet assignment problems (Hane, Barnhart, Johnson, Marsten, Nemhauser & Sigismondi (1995)), railroad car distribution problems (Jordan & Turnquist (1983), Mendiratta & Turnquist (1982), Haghani (1989), and Herren (1973), for example), the load matching problem of truckload trucking (Powell (1991), Powell (1996), Schrijver (1993)), routing and scheduling problems in less-than-truckload trucking (Powell (1986), Crainic & Roy (1991)), the flow management problem in air traffic control (Andreatta & Romanin-Jacur (1987) and Odoni (1986)) and the management of ocean containers (Crainic, Gendreau & Dejax (1993)). Even within an industry, rail car distribution is different from rail locomotive management or rail crew scheduling. Clearly, this tendency reduces our ability to learn from similar problems in different industries and excessively fragments the field.

This paper offers a new vocabulary for representing complex problems in a stochastic, dynamic setting. Our focus is on operational problems that need to be solved under uncertainty, and complex problems which are difficult to formulate mathematically in a clear and elegant way. We believe that similarities between problems are often disguised by semantic differences that reflect the contextual domain of an application. It is not readily apparent, for example, that the blocking problem of railroads and the load planning problem of less-than-truckload trucking are both instances of a dynamic service network design problem. The question that tends to arise is: When are apparently different problems similar, and when are seemingly similar problems different?

Consider, for example:

Example 1: Discrete scheduling

1a: A customer wants to move a truckload shipment from Cleveland to Atlanta. A motor carrier has a Dayton-based driver in nearby Dayton (creating the challenge of getting him back home after moving the load), and farther away an Atlanta-based driver sitting in Detroit (but the load would take him home). Which driver should the carrier use to move the load?

1b: A manufacturer sells speciality pipes cut to the length required by the customer. A customer order for a 12 foot length of pipe comes in. The manufacturer has bins of pipes ranging in length from two feet to 20 feet. Right now, he only has 15 and 20 foot sections of pipe. He can cut a 15 foot section, leaving a three foot remnant, or use a 20 foot section, leaving an eight foot remnant. Which length of pipe should he cut down to satisfy the customer order?

These apparently very different problems are actually almost identical, if viewed with the proper vocabulary. Both are instances of using heterogeneous, reusable resources to satisfy customer tasks. In each case, a resource is applied to a task which satisfies the customer, and produces a modified resource. In the case of the trucking example, the load modifies the driver resource by moving the driver to a new location. The pipe example modifies longer lengths of pipes and produces shorter ones. (This example is based on a conversation with Dan Adelman, who introduced a similar problem in fiber optics in Adelman, Nemhauser, Padron, Stubbs & Pandit (1996) in the context of a remnant inventory problem.)

In this paper, we offer a paradigm for representing a class of problems that we refer to as *dynamic resource transformation problems*. Since “DRTP” is relatively unpronounceable, we have chosen to refer to these problems as “DRiP’s” (if someone asks you what the “i” stands for, you are supposed to say “information”). Subclasses include *dynamic resource allocation problems*, which specifically address spatial problems, *dynamic resource scheduling problems*, which specifically address nonspatial problems (such as classical machine scheduling) or, more broadly *dynamic resource management problems* (by using the abbreviation “DRiP”, we invite the reader to choose any of the above descriptions that best fits their work). While much of what we cover will also apply to static problems, our focus is on stochastic, dynamic problems, which have received very little attention in the context of large-scale optimization problems. Our goal is to develop a vocabulary that is not only relatively context-free, but is also (again, relatively) free of a particular algorithmic strategy. In addition, our focus is primarily on very complex problems, such as those that arise in

freight transportation. These are problems that have classically been modeled as large scale linear or integer programs.

A number of papers have been published which offer representations of well-defined problems oriented toward identifying problems with common properties. The “ $M/G/k$ ” paradigm for queueing systems, introduced by Kendall (see Gross & Harris (1985) for a summary and references), served for decades to define the field of queueing theory. This framework actually takes the form “ $A/B/X/Y/Z$ ” where  $A$  and  $B$  capture the arrival and service processes, while  $X$ ,  $Y$  and  $Z$  represent the number of services, system capacity, and queue discipline. The machine scheduling literature uses “ $\alpha | \beta | \gamma$ ” to represent, respectively, the machine environment, processing characteristics, and the number of machines (see Pinedo (1995)). Eiselt, Laporte & Thisse (1993) propose a “I / II / III / IV / V” taxonomy for classifying location problems, and use this to organize the extensive literature in location theory. Bodin & Golden (1981) list 13 dimensions in their classification of vehicle routing problems, each of which is divided into several categories.

A different approach to problem representation has been modeling languages, which are more generally oriented toward much broader problem classes. An excellent review of this field is given in Geoffrion (1989*b*). Modeling languages can take two general forms: algebraic modeling systems, such as AMPL (Fourer, Gay & Kernighan (1990)) which are high level code or data generators for lower level systems, and conceptual representations, such as Geoffrion’s Structured Modeling Language (Geoffrion (1989*a*), Geoffrion (1989*b*), Geoffrion (1992*a*), Geoffrion (1992*b*)). The challenge in developing any representational system is the tradeoff between scope of application (generality) and ease of application (which requires domain-specific features). The more a representation is tied to a specific problem class, the easier it is to apply to that problem class, and the less useful it becomes to other problem classes. In our view, other work on general modeling approaches is difficult to apply to the stochastic, dynamic problems that arise in complex arenas such as freight transportation. This is not to say that they cannot be applied, but rather that we cannot easily determine how to apply them.

A major goal of the paper is an explicit representation of this problem class. This requires identifying the major elements of a DRiP in a way that could be represented in a computer dataset. As an illustration, a linear program, expressed as  $\min cx$  subject to  $Ax = b, x \leq u, x \geq 0$ , is expressed by specifying  $A, b, c, u$  and  $x$ . This representation has provided a highly flexible mathematical vocabulary, as well as serving as the basis for the venerable “MPS” format. In fact, one could argue

that Dantzig’s most valuable contribution was not the simplex algorithm, but the basic problem representation that has served as a useful language for a wide range of algorithms that do not even use the simplex method.

The central thesis of this paper is that we lack a basic representational paradigm for the types of stochastic, dynamic problems that arise in the context of complex operational applications such as freight transportation and logistics. These problems have been classically formulated as large-scale mixed integer optimization problems, where the primary emphasis is on finding an algorithm, while accepting strong simplifications on the representation of the problem. In this paper, we do not focus on an algorithm for solving a problem, but rather on simply representing the problem in a general way that captures the much richer set of modeling issues that arise in a dynamic setting. We feel that in a static, deterministic setting, the challenge of modeling the problem is relatively easy, while solving the problem is quite hard. In a stochastic, dynamic setting, the reverse is true. Since we will never obtain truly optimal solutions (for realistic problems), “solving” a stochastic, dynamic problem is relatively easy (since we do not generally look for optimal solutions), whereas simply modeling the problem can be extremely difficult. Of course, optimal solutions for this problem class are exceptionally hard, and the challenge of obtaining (and verifying) better solutions will remain an intellectually challenging task.

A goal of this paper is to provide a flexible representational structure that others can build on and adapt to fit their needs. Since our interest is in very complex problems, our intent is to build on existing modeling paradigms, rather than try to replace them with something completely new. At the same time, we feel that no single modeling paradigm, classically applied, allows us to capture the richness of complex stochastic, dynamic problems that arise in practice. The standard paradigm of most research in the operations research community is to propose a model, and then focus on developing an algorithm. In the case of large, complex problems, and in particular those arising in a dynamic setting, we feel that it is more important at this stage to agree on a mathematical representation that covers the important dimensions that distinguish a stochastic, dynamic problem.

The contributions of this paper are as follows: a) We define the DRiP as a problem class, and show how it includes a variety of problems in operations research. We provide a compact but *comprehensive* organization of the elements of a DRiP, which uses only three primary dimensions and ten subdimensions. We claim that this organization encompasses all the essential features of

this problem class, and as such, can become the basis for classifying problems (which we do not attempt). b) We provide a full notational system that is highly mnemonic and which captures the important elements of these problems in a natural way. Our mathematical representation provides easy interfaces with classical concepts in mathematical programming and stochastic processes. c) Our representation provides what we believe to be the first representation that focuses on the information content in decisions for a multiagent system, along with a novel characterization of information. Finally, d) we provide a structure for storing instances of a DRiP in files that can be shared.

The DRiP problem class offers several dimensions which appear to be new: 1) the concept of resource layering, whereby different types of resources can be combined to make composite resources; 2) the representation of complex physics through the use of a type of transfer function (called the *modify* function), that works at the level of each resource, with explicit treatment of the information content of each decision and the modeling of multi-period transfer times; 3) the modeling of problems that are typically handled as single, large optimization problems using the framework of multi-agent control; and 4) the modeling of informational subproblems, capturing the property in most large-scale systems that decisions are made in subproblems using subsets of information.

A challenge of this paper is striking a balance between specificity and generality. Individual researchers may react to our representation by claiming that it “forces” a lot of “unnecessary” complexity. Our intent in this paper is not to encourage any particular representation (we even provide for the explicit use of deterministic approximations). However, we do wish to highlight simplifying assumptions when they are made, so that perhaps other researchers may in time investigate the impact of these assumptions. Ultimately, the goal of modeling is finding the set of approximations that balances accuracy with tractability.

Section 1 gives a number of examples of these problems, illustrating the fundamental elements of a DRiP. Section 2 outlines the three primary and the associated secondary dimensions of a DRiP. The remainder of the paper focuses on the mathematical representation of this problem. Given the complexity of our problem class, we present our notational conventions in section 3. Then, sections 4, 5 and 6 give mathematical descriptions of the three primary dimensions of a DRiP. Section 7 shows how a DRiP can be represented as a set of files giving in the form of numeric data and software. Finally, section 8 summarizes the paper.

# 1 Examples

Dynamic resource transformation problems arise in a variety of settings. Each of these involves the transformation of one or more types of resources in a way that not only generates costs, but also some sort of benefit that justifies the cost of transformation. Although our interest is primarily in transportation-related problems, we include nontransportation examples to stress the generality of the problem class.

A central feature of our problem class is that resources often occur in layers. Problems with more layers are inherently more difficult to solve than problems with fewer layers. Below, we list a series of applications, organized by the number of layers in each example.

- One layer problems:
  - Classical inventory planning (but with no demand backlogging).
  - Distribution problems (possibly multiproduct) with stochastic demands (but no demand backlogging).
  - The dynamic travelling repairman problem.
  - Machine scheduling (with no setups).
  - Airline crew scheduling problems (with fixed times for flights).
  - Truckload fleet management I: managing the flows of (perhaps multiple types of) trucks over a set of loads with fixed departure and arrival times.
  - Less-than-truckload trucking I: the traffic assignment problem (routing shipments).
  - Rail operations I: distributing box cars to customers with no demand backlogging.
- Two layer problems:
  - Vehicle routing problems (vehicles and product).
  - Machine scheduling with setups (jobs and machines).
  - Personnel planning (assigning people to jobs).
  - Fleet management (vehicles and loads or customers to be moved).
  - Truckload fleet management II: optimizing the assignment of drivers to loads, both of which may be served over specified periods of time (time windows).

- Less-than-truckload trucking II: routing shipments and trailers.
- Rail operations II: Scheduling locomotives to move trains.
- Three layer problems:
  - Truckload fleet management III: (simultaneously managing drivers, tractors and trailers).
  - Machine scheduling (with setups) with operators (jobs, machines and people).
  - Less-than-truckload trucking III: routing shipments, trailers and drivers.
  - Rail operations III: Scheduling locomotives and crews to move trains.
- Four layer problems:
  - Routing and scheduling for chemical distribution (driver, tractor, trailer and product).
  - Multiple machine scheduling (with setups) with operators (jobs, two types of machines and people).
  - Less-than-truckload trucking IV: routing shipments, tractors, trailers and drivers.
  - Rail operations IV: Scheduling locomotives, crews and boxcars to satisfy external customer demands that may be backlogged.

DRiP's involve objects with attributes that evolve over time as a result of various physical processes. Examples of these processes include:

- Temporal processes - These are physical processes that govern how long a task or operation takes. These include travel times, time required to execute a task, time to clean a trailer, time to train a person, and off-duty and sleeping time.
- Economic processes - Each transformation is assumed to generate one or more measurable quantities that are accumulated directly in an objective to be optimized. These include the generation of costs, revenues, service measurements and measures of equipment productivity or employee satisfaction.
- Discrete classification processes - Resources often have one or more discrete classifications that can change. This might be the location of a truck, the state of a machine, or the skill level of a consultant.

- Aging and replenishment processes - Many resources have a dimension that undergoes aging and replenishment. Examples include consumption of allowable work hours, fuel consumption and deterioration of equipment.
- Arrival and departure processes - These describe the arrival and departure of resources to and from the system.
- Information processes - Most important, information processes describe the arrival of data into the system that can be used for making decisions.

The last dimension of a DRiP is the means by which the system can be controlled. Most complex systems are run by more than one decision-maker, referred in the literature to multi-agent control, while classical optimization models assume a single-agent structure. Given a control structure, we need then to characterize the means by which decisions are made. A decision function may be rule-based (given a state, choose an action) or based on some form of cost minimization. Simple problems may use cost, profit or level of service, but more complex systems combine multiple measures, and may not even have quantifiable evaluation criteria.

## 2 Dynamic resource transformation problems

This section provides a conceptual overview of a DRiP, summarizing all the major dimensions and subdimensions, and discussing some of the conceptual issues. This section is then followed by a more formal mathematical representation of all the elements in sections (4), (5) and (6).

The first step in the development of a representational system is to describe the dimensions of a dynamic resource transformation problem. A DRiP, in our view, is comprised of three fundamental elements:

- Resources - These are the set of data elements organized into classes which make up our system. The term resources can be used narrowly to mean physical resources (drivers, aircraft, components) or more broadly (financial resources, informational resources).
- Processes - Here we capture the physics of the problem, comprised of the laws that govern how the system evolves over time, physical constraints that must be observed, and the evolution of information coming to the system.

- Controls - Controls describe what decisions need to be made, who makes them, how information is provided to the decision maker, and how the decision maker separates good from bad decisions.

We use this fundamental organization to represent an instance of a DRiP. We use the general style of queueing theory and machine scheduling, and propose that dynamic resource transformation problems be identified using the notation:

$$Information \parallel Processes \parallel Controls$$

where each element (resource, process or control) can be replaced by a series of descriptive abbreviations that describe the nature of that element. As such, our organization can be used as a basis of a taxonomy. For reasons of space, we do not attempt to actually design such a taxonomy in this paper.

As this paper was being written, there was active consideration of whether the first dimension should be called “resources” (which for many people refers to physical objects), “information” (to more broadly capture the set of data that may or may not encompass physical resources) or “knowledge” (which brings into play subtle distinctions between “information” and “knowledge”). If the reader is considering a problem involving the management of physical resources (drivers, trucks, planes, inventory), then the only interesting type of information might be the resources themselves. In this case, we invite the reader to use the notation “*Resources*||*Processes*||*Controls*”. If it is important to explicitly model information that is not, strictly speaking, a resource (for example, the travel times on a network or the price of fuel), then we recommend retaining the “information” dimension, and explicitly modeling the larger information set.

Some readers may argue that our first dimension, strictly speaking is “knowledge.” We do not disagree with this assessment, but feel that the use of the term “information”, while lacking precision, is not incorrect (we view knowledge as a subset of information).

We now summarize the subelements of each one of these dimensions. Our goal in this section is to provide a *comprehensive* summary of the elements of a DRiP. We claim that our list of three primary dimensions, and the nine subdimensions given below, is *complete* and comprise all the dimensions needed to express this problem class. We acknowledge that different people might organize the information within the subdimensions differently (in the course of writing this paper,

we progressed through a range of organizations before settling on the structure presented here).

In our structure, we identify the subdimensions of a problem using numbers, as in I.1 (for the first dimension of information) or P.1.a (as in subelement 1.a of processes). Section 2.1 outlines the elements of information, section 2.2 does the same for processes, and section 2.3 summarizes the elements of controls.

## 2.1 Information

The information available to our system may be organized as follows:

- I.1) The information classes and attributes - Here we list the different classes of information that make up our problem. This includes the attributes of each information class, and the subsets that make up the resources that are being managed.
- I.2) Aggregation functions - Aggregation is the means by which we express the information relevant to a decision. It is sometimes necessary to use multiple levels of aggregation within the same model; for this reason, the aggregation function (or a family of aggregation functions) must be an explicit part of the modeling paradigm.

Information can be organized as follows:

- 1) Information describing the state of the system. This can be organized into two groups:
  - a) Information about resources to be managed.
  - b) Information that describes the physics of the system.
- 2) Plans or policies which provide an indication of how the system is going to evolve over time. A plan may be a set of decisions at some level of aggregation that provide guidance regarding future decisions.

Most optimization models are formulated without the presence of an exogenous plan. Unfortunately, this creates a number of practical difficulties, which is the reason why virtually all large organizations develop plans and use these as a rough guideline.

We are particularly interested in information describing resources. It can be argued that there are three types of resources: physical, financial and informational. There are many applications

(routing and scheduling, physical distribution, fleet management, personnel management) where the only resource is a physical resource. It is common in such settings to drop the adjective “physical” because these are the only resources being managed. For the most part, this paper focuses on physical resources, and this should be assumed unless it is specifically indicated otherwise. However, we invite the reader to consider other applications.

Resource classes can be broadly categorized as follows:

- Static classes are classes with data that will never change within the context of a model.
- Dynamic classes have data that will change. This class can be further categorized by *how* the attributes of the object change. Attributes may change as a result of:
  - Exogenous processes, such as weather, equipment failures, or decisions made by exogenous agents, or
  - Endogenous processes, reflecting decisions made within our system.

The concept of a resource is especially important in our paradigm since our focus is on how to control this class of elements. For this reason, we need a formal definition of a resource:

**Definition 2.1** *A resource is an endogenously controllable information class which constrains the system over time.*

A grey area in the definition of a resource is the modifier “over time.” It may be that something constrains the system only at a point in time. Such constraints can certainly be considered resources, but we later argue that this is optional (they can be captured as a simple form of constraint, which we view simply as a model parameter). By contrast, if an object constrains a system over multiple time periods (during which new information may be arriving), then it generally *must* be modeled as a resource class.

It is important to distinguish between two types of resources:

**Definition 2.2** *An active resource class is a resource with attributes that can be endogenously modified without necessarily being coupled with other resources. A passive resource layer is one that can only be coupled or uncoupled with other resources.*

Active resources are the resources that we are managing. Passive resources enable the active resources to do their job (and if there are not enough of them, then they prevent active resources from doing their job).

The concept of active and passive resources is more precisely defined in sections 2.3 and 6 when we formally define what we mean by controls. In the case of trucking, we might consider three layers of resources: loads, drivers and trucks. A truck and a driver may move empty, and a load, truck and driver may represent a loaded move. Assume that a truck by itself cannot do anything. In this case, the truck would be a passive resource, while the driver with the truck can move empty.

**Definition 2.3** *A **persistent** resource is a resource that remains in the system over the planning horizon (the resource may not be present in the system initially, but remains in the system once it arrives). A **transient** resource may enter and leave the system during the planning horizon (some fields refer to these as **perishable** resources).*

Important classes of persistent resources include:

**Definition 2.4** *A **recurrent** resource is a persistent resource that needs to continually cycle back to a particular base state. A **strongly recurrent** resource needs to return to a base state in a fixed period of time (drivers returning home at the end of the day; aircraft that have to cycle back for a legally mandated maintenance check). **Weakly recurrent** resources have an incentive to cycle back, but the force of this requirement is not as strong. **Nonrecurrent** resources, such as trucks, trailers and containers, often move aimlessly around the system with no particular constraint to return to a particular state.*

Most applications involve a resource layer that can be viewed as a task, customer, job or demand. If a customer demand may be satisfied over some interval of time (a time window) then it satisfies our strictest definition of a (passive, transient) resource: it is an information object, it is endogenously controllable (since we determine when within the window it should be satisfied) and it constrains the system (over time). We acknowledge that many readers will probably have initial difficulty with the notion that a “demand” is a type of “resource.” We do not feel that there is any difficulty labeling a resource class as a task or demand, but notationally, it is important to understand that tasks are just a special case of a resource.

An important distinguishing characteristic of transient resources is the set of entry and exit points from the system. We propose the following definitions:

**Definition 2.5** **Point** resources have a single point of entry and exit (e.g. a customer demand at a point in the system which may be backlogged). **Path-based** resources have a single point of entry and exit which are different. **Tree-based** resources have a single point of entry with multiple exit points or a single point of exit with multiple entry points. **General** transient resources have multiple entry and exit points (e.g. agricultural commodities that can be produced and consumed at multiple points in the network).

An important dimension of DRiP's is the ability to build up complex resources from simple ones. We define:

**Definition 2.6** A **primitive** resource is an elementary, indivisible resource with a fixed set of attribute types and predefined behavior. A **composite** (or **compound**) resource is formed by joining two or more resources in different classes to make a resource with a combined set of attributes.

**Definition 2.7** We **bundle** two resources by joining two resources in the same object class. We **couple** two resources in different classes. We bundle primitive resources and couple them together to form composite resources.

A composite resource is a resource with attributes drawn from multiple resource *classes*. For this reason we define:

**Definition 2.8** A **resource layer** is a set of attributes drawn from one or more resource classes.

Resource layers would be defined based on natural couplings of resources, such as “pilot  $\times$  aircraft”. While these are largely determined by the physical realities of the problem, the choice of resource layers will generally be up to the modeler. It is useful to define:

**Definition 2.9** A **primitive layer** is a resource layer with attributes drawn from a single object class; a **composite layer** has attributes drawn from more than one class.

The concept of layering arises frequently in complex operations. Not all problems need the complexity of layering (many problems can be effectively modeled as one-layer problems), but it is

important to provide the framework to handle more complex problems, where multiple types of resources may be combined to form a new resource with a richer set of behaviors.

Compound resources arise in many settings. When the resources are people, they are called teams. If the resources are financial assets, we call compound resources portfolios. When the resources are components being assembled, we refer to the compound resources as a product (to be sold).

Sometimes it is a bit tricky defining the set of resources, keeping in mind that we want to define the simplest set of resources possible. In machine scheduling (in particular, problems involving machines with setups), there would be two resource layers: machines and jobs. If the machines do not have setups, then we can model this problem with one resource layer, consisting only of jobs. The operations research literature has long used the vocabulary of managing resources (such as people and machines) to serve customers (or tasks). In our view, tasks and jobs are instances of *customer orders* which is a form of resource. We make a clear distinction between *customer orders*, which are endogenous “resources” to be managed, and *customers*, which are typically exogenous to the system.

## 2.2 Processes

The elements of processes are:

- P.1) Information processes - Systems are affected by information processes that impact the transformation of a resource. There are two types of information processes:
  - P.1.a) Exogenous information processes, which represent information updates from outside the system.
  - P.1.b) Endogenous information processes, which are the sequence of decisions made within the system.
- P.2) System dynamics, which are the physical laws that govern the evolution of the system over time. There are three subelements in this component:
  - P.2.a) The modification of attributes - These are the equations governing how the attributes of a resource change over time in response to endogenous and exogenous controls.
  - P.2.b) The economics of a transformation (generation of costs and rewards).

P.2.c) The time required to effect a transformation, called the *transfer time*.

P.3) Constraints on transformations - These restrict our ability to transform a process. We provide for exactly two classes of constraints:

P.3.a) Conservation of mass (a resource can not be in two places at the same time; you cannot create or destroy resources, although they can enter and leave the system).

P.3.b) Rate of process transformation - Here we explicitly account for the presence of multiple resources where there is a restriction on the number of resources that can be modified at any given time. Specific classes of rate transformation constraints include:

P.3.b.i) Technology constraints - The speed of a machine or the size of a truck both represent technological parameters that restrict the rate at which a resource may be transformed.

P.3.b.ii) Exogenous controls - The rate of process transformation may be limited by exogenous factors such as policies set by higher levels of management.

P.3.b.iii) Market demands - We provide a special category to represent the effect of market demands, since satisfying a customer enables a special type of resource transformation, that is limited by the size of the market. However, a market demand can be represented by a constraint only in very special cases (no advance information, and no demand backlogging).

The concept of decisions as an information process is unusual, but we will see later that exogenous information and decisions both play a similar role. Furthermore, this characterization will lead us to a nonstandard mathematical description of the problem.

Virtually all of the physics of a real problem are contained in the dimension called system dynamics. It is common in classical mathematical programming paradigms to mix system dynamics with what we call constraints into a group of equations called constraints. What we refer to as a constraint is much more limited than what is normally modeled as a constraint in an optimization model. For example, the time required to effect a transformation is a type of constraint, but we capture this in the transfer function. Similarly, time window constraints may be modeled in the cost function (e.g. a high cost for transformations that are outside the time window) or through rules that govern allowable transformations.

## 2.3 Controls

There are three broad classes of controls:

**Primal** - These are decisions that either directly change the state of a resource, or the parameters that determine the outcome of a decision.

**Dual** - These are parameters that affect the costs (or rewards) generated by a decision.

**Informational** - In special cases, we can choose to provide additional information to the process of making a decision, as might occur when we forecast future events and make plans in the future.

It is important to recognize the presence of all three classes of controls, but our focus will be on primal controls. We define two broad classes of primal controls:

**Definition 2.10** **Direct controls** are the only means by which the attributes of a resource may be changed. **Indirect controls** are parameters that can be changed endogenously and which have an effect on the outcome of a direct control.

Primal, direct controls are what we typically think of as decision variables: moving the truck, making product, managing inventories. In dynamic systems, these decisions are invariably state dependent, and need to be made as the system evolves. Indirect controls represent a kind of “everything else” category that covers parameters that impact the results of a decision (the speed of an aircraft, the capacity of a roadway).

For the remainder of our presentation, we focus only on a discussion of (primal) direct controls. Recall that primal, direct controls are a type of information process since they are generated as the system evolves. As a result, the definition of these controls is given under subdimension P.1.b of processes. The controls dimension of our representation, then, focuses on how decisions are made. There are four subdimensions of this component:

- C.1) The control structure - Since large problems are often divided into zones of control (often called multiagent systems) we have to specify who owns what decision.
- C.2) The information set - An important dimension of complex problems is specifying what information is available when a decision is made.

C.3) The decision function - This covers how decisions are made, and with what information.

C.4) Measurement and evaluation - Finally, we have to specify how we compare one decision to another.

Capturing the set of direct decisions is helped by defining:

**Definition 2.11** *A primitive decision is an elementary action that modifies the system and which cannot be represented as a sequence of other decisions.*

**Definition 2.12** *A composite decision or tactic is a sequence of two or more decisions that can be executed with the same set of information.*

A tactic might be a simple sequence of actions. For example, a cab company might have a tactic: “serve a customer” which consists of: 1) move from current location to the pickup point of the customer (a form of modify called a “move”), 2) pick up the customer (a coupling), 3) go to the customer’s desired destination (another “move”), and 4) drop off the customer (uncouple). It is not uncommon to predefine certain classes of tactics. When this is the case, a tactic is handled just like a primitive decision, but it may impact more than one resource (or type of resource).

There are three classes of primitive controls. These are:

Couple - Putting resources in different layers together, such as putting a pilot in an aircraft.

Uncouple - Such as taking the pilot out of the aircraft.

Modify - For example, using the pilot to fly the aircraft to another location, producing a change in location, fuel, and pilot hours.

We claim that the three fundamental decisions (couple, uncouple, and modify) are the *only* ways a DRiP can be directly controlled (endogenously), and that all decision variables are instances of this subset. The challenge in a DRiP is to choose a set of controls over time to achieve specific objectives, represented in the evaluation dimension.

Important special cases of modify include:

1. Hold (do nothing).

2. Move (a resource from one spatial location to another).
3. Entry (allow a resource to enter the system).
4. Exit (allow a resource to leave the system, a process).

The option to “hold” a resource is important, since this “decision” is generally implicit in the lack of information to the contrary. Notationally, we provide the following special notation for *hold* decisions:

$d^\phi$  = The decision to “hold” a resource.

Later, we provide specific conditions for the presence of the “hold” decision.

### 3 Notational style

Since this is a modeling paper, precise, elegant notation is critical. For this reason, we follow a number of notational conventions that are designed to simplify the process of modeling complex operations. If we have a vector  $a$ , we denote elements of the vector at all times using subscripts. Since we have many quantities defined over time, we may view  $a$  as the vector of all elements over all points in time, and  $a_t$  as the vector of elements at time  $t$ , with individual element  $a_{it}$  or  $a_{ijt}$ . We often need different flavors of a variable, such as different types of costs, or different types of decisions. Rather than using different variable names, we generally use the same variable with a superscript to denote a particular flavor. Thus, we might use  $c^h$  as a holding cost and  $c^f$  as a fixed cost. If one flavor is used most of the time, we may define this without a superscript.

We often need to update a variable iteratively. It is important to use a common variable such as  $n$  ( $k$  is also often used) to represent the iteration counter.  $n$  is useful as an iteration counter because variables such as  $i$ ,  $j$ ,  $k$  or  $\ell$  often have other uses. Also, some algorithms have inner and outer iterations; in this case, we recommend using  $n$  as the outer iteration counter, and  $m$  as the inner iteration counter. We always put the iteration index in the superscript (as in  $x^n$ ). If we are using the superscript to represent the flavor of a variable, we may need to use a double superscript, as in  $x^{\ell,n}$ ; in this case, the iteration counter ( $n$ ) *must* be the outer index (since “ $x^\ell$ ” is the name of the variable). It is sometimes necessary to have different flavors of the *same* variable. In this case, we recommend using notation such as  $\hat{x}$  and  $\bar{x}$ .

We make extensive use of sets. Sets are denoted using the calligraphic font  $\mathcal{A}$ , and subsets are denoted using a superscript in parentheses, as in  $\mathcal{A}^{(c)}$ . Throughout, we use the word set in the sense that it represents a space of possible outcomes of some variable or function, and not simply a collection of elements (as in a vector). Thus, for any set such as  $\mathcal{A}$ , we would later make a statement such as  $a \in \mathcal{A}$ . Following our convention on superscripts,  $\mathcal{A}$  represents a set, while  $\mathcal{A}^s$  represents a different set;  $\mathcal{A}^{(c)}$  represents a subset of  $\mathcal{A}$  (similarly  $\mathcal{A}^{s(c)}$  would be a subset of the set  $\mathcal{A}^s$ ). We always use specific lowercase letters to index a set (almost always the lowercase version of the set), as in  $a \in \mathcal{A}, t \in \mathcal{T}, c \in \mathcal{C}$ . Thus,  $x_t$  always refers to  $t \in \mathcal{T}$ , while  $x_i$  would always refer to  $i \in \mathcal{I}$ . The element  $x_{it}$ , then, refers to the element  $i \in \mathcal{I}, t \in \mathcal{T}$ . If we need to simultaneously refer to two different elements of a set, we would use  $a, a', a'', \bar{a}$  or  $\hat{a}$  to refer to elements of  $\mathcal{A}$ . The letter  $b$  will *never* be used as an element of  $\mathcal{A}$ . We say that we *almost* always use the lower case version of a set to index the set, because some exceptions are bound to arise. For example, it is *very* common in transportation to index a location by  $i$  and  $j$  (as in  $x_{ij}$ ; our notation would require us to say  $x_{ii'}$  or even  $x_{\ell, \ell'}$  for locations  $\ell \in \mathcal{L}$ ). We feel it is acceptable to define  $i$  and  $j$  to be locations in, say, the set  $\mathcal{I}$  (or some other set), but would insist that  $i$  and  $j$  *always* refer to a city.

Every effort has been made to make the notation as mnemonic as possible. For example,  $x$  is *always* a decision variable,  $t$  always refers to time, and so on. When we need different variables for time, we use  $t', t''$  and so on. We make extensive use of sets, and stringently avoid numerical indexes such as  $x_1, x_2, \dots$ . We have worked to make our notation “friendly” to software implementation, and have adopted several C programming language conventions. The extensive use of sets is important, since we often use subscripts which are themselves vectors. Thus,  $x_a$  might be the number of elements with attribute vector  $a$ . Summing over these elements, then, is accomplished using notation such as  $\sum_{a \in \mathcal{A}} x_a$ . We feel that this is preferable to creating an index  $i$  where  $\mathcal{I} = \{0, 1, \dots, |\mathcal{A}|\}$  are the elements in the set since the numerical indexes have no inherent meaning.

Because of the importance of stochastic processes in our modeling paradigm, we need to standardize our notation for random variables. It is fairly standard to represent a probability space using notation such as  $(\Omega, \mathcal{F}, \mathcal{P})$  where  $\Omega$  is a set of elementary outcomes,  $\mathcal{F}$  is a set of events (think of it as a collection of the outcomes of a set of random variables), and  $\mathcal{P}$  is a probability measure defined on the space  $(\Omega, \mathcal{F})$  (giving us, in effect, the probability of an outcome in the set  $\mathcal{F}$ ). The notation  $\Omega$  and  $\mathcal{P}$  is fairly standard, but it is not unusual to see  $\mathcal{E}$  (set of events)  $\mathcal{H}$  (set of histories) used instead of  $\mathcal{F}$ . When information arrives over time, it is common to define a series of

subsets  $\mathcal{F}_t$  (or  $\mathcal{H}_t$ ) which capture what is known by a particular point in time (if  $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$ , then the process is a filtration, which is the reason behind the popularity of the notation  $\mathcal{F}$ ).

One of the most difficult challenges is in the representation of a random variable, partly because there is not a single standard in the research literature. Some common conventions for representing random variables (capital letters, bold letters, letters with hats, the letters  $X, Y$  or  $Z$ ) are not generally workable in engineering applications. One easy way to express that a function  $f$  is a random variable is to write it using functional notation:  $f(x, \omega)$ , which indicates that the function depends on both  $x$  and  $\omega$ . This is fine, but the reader needs to realize that the probability community will often write a function, say  $S_t$  (the state of the system at time  $t$ ) without indicating what the function depends on (the functional notation  $S_t(x, \omega)$  is more common in engineering). It is common to simply define a function, say  $f_t$  to be, say, a  $\mathcal{F}_t$ -measurable function (which means that  $f_t$  is a function of the events in  $\mathcal{F}_t$ ), and then the reader simply has to remember that the function is random (it is not indicated in the notation). This is rarely an issue until it is necessary to take an expectation. Thus, if  $f_t$  is random and  $g_t$  is not, then  $E\{f_t + g_t\} = Ef_t + g_t$ . In some communities, it is perfectly acceptable to write  $E\{f_t(x, \omega) + g_t(x)\} = Ef_t(x, \omega) + g_t(x)$ ; the reader must be warned, however, that this is not acceptable in the applied probability community, which views  $\omega$  as a number (not a random variable, which is always viewed as a function of  $\omega$ ), which means that  $Ef_t$  is the expectation of the random function  $f_t$ , while  $Ef_t(x, \omega) = f_t(x, \omega)$ , since given  $\omega$ ,  $f_t(x, \omega)$  is no longer random.

## 4 Information

The information dimension of our representation captures what we know about our system. In this sense, the dimension is somewhat misclassified (it could easily have been called “knowledge”). There are two dimensions to our information dimension: a) the information classes and their attributes, and b) aggregation functions, which capture what is important.

### 4.1 The information classes and attributes

Information can be divided between data describing the state of our system, and (optionally) a plan, that provides, at some level of aggregation, a summary of decisions which we anticipate will be made in the future. The concept of a plan is often (but not always) ignored in the formulation

of optimization models; because it is usually critical in real world operations of complex systems, we include it here for completeness. We defer the representation of a plan until later.

To capture the state of the system, we begin by defining:

$\mathcal{C}^I =$  Set of information objects that organize all the data in our system.

$\mathcal{E}_t^{(c)} =$  Set of information elements within class  $c \in \mathcal{C}^I$  at time  $t$ .

Each information element has an attribute vector given by:

$a_e =$  Vector of attributes of element  $e \in \mathcal{E}$ .

$\mathcal{A}^{(c)} =$  Attribute space of elements in class  $c$ , where  $a_e \in \mathcal{A}^{(c)}$  for  $e \in \mathcal{E}^{(c)}$ .

We capture all the information in our system using:

$K_t =$  The knowledge base (equivalently, the state of the database) at time  $t$ .

$= \{a_e, e \in \mathcal{E}_t^{(c)}, c \in \mathcal{C}^I\}$

The set  $K_0$  is given to us as data. Later, we formalize how  $\mathcal{E}_t$  and  $K_t$  evolve over time.

We need to divide information classes based on how the attribute vector  $a_e$  evolves. Let:

$\mathcal{C}^{I,s} =$  Set of information classes with static attributes.

$\mathcal{C}^{I,d} =$  Set of information classes with dynamic attributes.

We now divide the classes with dynamic attributes into two groups: those with attributes that change purely as a result of exogenous events, and those that evolve at least in part as a result of endogenous events. Following our earlier definition, we create a special class of elements which fall in the category of resources:

$\mathcal{C}^R =$  Set of information objects that comprise our *resources*.

$\subseteq \mathcal{C}^{I,d}$ .

The basic data for resources is captured using:

$\mathcal{R}_t^{(c)}$  = Set of individual resources (in class  $c$ ) at time  $t$ . This is the same as  $\mathcal{E}^{(c)}$  for  $c \in \mathcal{C}^R$ , but we need mnemonic notation that specifically refers to the set of resource elements.

$a_r$  = Attribute vector of the  $r^{th}$  resource,  $r \in \mathcal{R}_t$ .

$\mathcal{A}^{(c)}$  = Space of possible attributes of resources in class  $c$ .

We can think of  $a_r$  as data that would be read in from a file for a given resource  $r$ , and  $\mathcal{A}$  is the set of possible values of each element of  $a$ . A crew scheduling application, for example, might have an attribute vector  $a = \{a_{location}, a_{DOThours}, a_{skillset}\}$ . Then,  $\mathcal{A}_{location}^{(c)}$  would be the set of possible locations for an element in class  $c$ , and  $a_{location,r}$  would be the location of the  $r^{th}$  driver (at a particular point in time).

An important issue arises when representing the difference between what we know about the system now, versus what we think will be true about the future. Thus, a vector of attributes  $a_r$  may describe the resource not as it exists “here and now” but rather as we think it will exist at some point later (for example, a resource  $r$  is moving from  $i$  to  $j$ , and is currently enroute between these locations; the vector  $a_r$  may describe the attributes of the resource as we think they will exist when it arrives at  $j$ ). For this purpose, we define:

$a_{time}$  = The time at which the attribute vector  $a$  will become effective.

If numeric indices are being used (as in  $a_0, a_1, \dots$ ), then we suggest that  $a_0$  be used to capture the time at which the elements  $a_j, j > 0$  will become effective. If  $a_{time} > t$  (where  $t$  is the current time), then the vector  $a$  is likely to be random, since new information may still arrive. We use the notation:

$$a_i = \phi$$

if we do not have any information about the  $i^{th}$  element. Allowing an attribute vector  $a$  to take on a set of “estimates” which may yet change as time evolves reflects, we believe, common practice.

The remainder of this section provides discussions of two special types of resources. Section 4.1.1 discusses simple classes of resources that are more commonly referred to as “commodities” in the operations research literature. Then, section 4.1.2 introduces a very complex set of *composite*

*resources* create by combining the attributes of different types of primitive resources. Finally, section 4.1.3 reviews some important concepts in the definition of state spaces.

#### 4.1.1 Multiattribute resources and commodities

We wish to formalize the relationship between a multiattribute resource and what are often referred to as multicommodity flow problems. To do this, we begin by observing that the elements of the attribute vector  $a$  can be divided into two broad groups of elements:

$a = (a^{(s)}, a^{(d)})$ , where:

$a^{(s)}$  = The static elements of  $a$  that do not change over time.

$a^{(d)}$  = The dynamic elements of  $a$  that do change over time.

We then define:

$\mathcal{A}^s$  = The space of possible outcomes of  $a^{(s)}$  with element  $a^s$ .

$\mathcal{A}^d$  = The space of possible outcomes of  $a^{(d)}$  with element  $a^d$ .

For example, we might represent the attributes of a generic driver for a trucking company by the attributes:  $a = (a_{location}, a_{domicile}, a_{bid\_status}, a_{sleep\_status}, a_{driving\_hours}, a_{duty\_hours})$ . Thus,  $a^{(s)} = \{a_{domicile}, a_{bid\_status}, a_{sleep\_status}\}$  would represent the static elements, while  $a^{(d)} = \{a_{location}, a_{driving\_hours}, a_{duty\_hours}\}$  would be the dynamic elements.

We make this observation because it is common in some communities to view different types of resources as commodities. To complete this important relationship, let the elements of  $\mathcal{A}^s$  be indexed by  $k$ , and let:

$$\mathcal{K} = \{1, 2, \dots, |\mathcal{A}^s|\}$$

Then we may write:

$$\mathcal{A}^s = \{a_k^s\}_{k \in \mathcal{K}}$$

Next, we define a set of *states*:

$$\mathcal{I} = \{1, 2, \dots, |\mathcal{A}^d|\}$$

$$\mathcal{A}^d = \{a_i^d\}_{i \in \mathcal{I}}$$

(The use of the non-mnemonic labeling  $i \in \mathcal{I}$  for states follows the standard convention for multi-commodity flow problems; it also avoids confusion between “static” and “state.”)

This notation allows us to provide a definition of the concept of a commodity that is broader, and more formal, than is typically used in the literature:

**Definition 4.1** *A commodity  $k$  is a set of resources  $\mathcal{R}^{(k)}$  defined by:*

$$\mathcal{R}^{(k)} = \{r | a_r^{(s)} = a_k^s \in \mathcal{A}^s\}$$

If the concept of a commodity gives us a short-hand notation for the static attributes of a resource, then we need a short-hand notation for the dynamic attributes of a resource. For this purpose, we might define:

**Definition 4.2** *The state  $i$  of a commodity  $k$  is the vector of dynamic attributes  $a_r^{(d)} = a_i^d$ .*

Thus, if  $a_r$  is the attribute vector of resource  $r$ , then we would say that resource  $r$  is a commodity of type  $k$  in state  $i$  if  $a_r^{(d)} = a_i^d, r \in \mathcal{R}^{(k)}, i \in \mathcal{I}$ .

In some applications, the dimensionality of the vector  $\mathcal{A}^{(d)}$  is quite small, representing a type of product (wheat or corn), a type of vehicle (Boeing 747, Boeing 727) or a pilot (identified by a unique ID). In these problems, the concept of a “commodity” is useful. However, the single word “commodity” is awkward when there are different resource layers, creating, for example, different commodities that represent not only different types of aircraft, but also different pilots, as well as the combination of aircraft and pilots. If a reader is comfortable with the term commodity, we would suggest introducing the concept of *commodity classes* so that different classes of commodity (aircraft and pilots) can be distinguished from different types of resources within a commodity class.

### 4.1.2 Composite resources

One of the most challenging dimensions of more complex DRiP's is the representation of different types of resources and their interactions. The bundling, coupling and uncoupling of resources arises because in many problems, resources can be organized in *layers* which can be coupled together to acquire new behaviors. In some cases, resources in the same layer may work together, such as arises when multiple locomotives are joined to create a single power unit (called a consist) to pull a train.

It is important to distinguish between resource classes,  $\mathcal{C}^R$ , and resource layers. Resource classes might be a pilot, aircraft, navigator or fuel. A resource layer is generally a resource coupled with one or more other classes (or layers). Resource layers arise when we want to define resources that may be coupled with other resources. Let:

$\mathcal{L}$  = Set of resource layers (each comprised of a set of one or more resources classes coupled together).

If a layer consists of resources from, say, classes 1, 2 and 3, then the attribute vector of this layer is given by:

$$\begin{aligned} a^{1|2|3} &= \text{The attribute vector of layer 1} \\ &= a^1|a^2|a^3 \end{aligned}$$

We illustrate our approach to layering using an air freight example with four resource classes: (1) fuel, (2) aircraft, (3) pilot and (4) freight. Let the attribute vector of each class  $i$  be given by  $a^i \in \mathcal{A}^{(i)}$ . There is a natural hierarchy in these layers:  $1 \rightarrow 2 \rightarrow 3 \rightarrow 4$ . That is, we put fuel in the aircraft, then add a pilot, and then add freight. The combination fuel|aircraft|pilot is an active composite resource (for example, I can fly this aircraft empty from one city to another), or I can add in freight, and fly the aircraft loaded. We represent these possible combinations using:

$$\mathcal{L} = \{1, (2|1), (3|2|1), (4|3|2|1)\}$$

The first layer is the fuel layer, which is a primitive resource. The second is the aircraft layer, which includes the aircraft and fuel, and so on. We label each layer by its lead resource, so the pilot layer

is comprised of pilot|aircraft|fuel. We represent the attribute vector of each layer using:

$$\begin{aligned} a^{3|2|1} &= a^3|a^2|1 \\ &= a^3|a^2|a^1 \\ &\in \mathcal{A}^{(3|2|1)} = \mathcal{A}^{(3)} \times \mathcal{A}^{(2)} \times \mathcal{A}^{(1)} \end{aligned}$$

The vector  $a^{3|2|1}$  is formed of the original attribute vector  $a^3$  with additional elements that represent fields indicating the potential for coupling with layer 2. Note that this field may be empty, indicating (in this example) a pilot with no aircraft. Thus, a pilot layer may consist of just a pilot, or a pilot coupled with an aircraft (which may in turn be coupled with fuel). Clearly, the field that captures which aircraft a pilot is coupled with is a decision variable and needs to be determined endogenously.

With these concepts in mind, we define:

**Definition 4.3** *A layering is a set  $\mathcal{L}$  of subsets of resource classes  $c \in \mathcal{C}^R$ , representing all the possible composite resources that may arise in practice. A layer consists of a **lead resource class** and any resource classes that may be combined with with the lead resource class. A resource class may be the lead resource class for only one layer.*

Typically, each layer is given the name of the lead resource class, meaning that the number of layers is equal to the number of resource classes.

We refer to attributes of layers using:

$$b^\ell = \text{Attributes of layer } \ell \in \mathcal{L}.$$

$$\mathcal{B}^{(\ell)} = \text{The space of attributes of layer } \ell.$$

The “ $b$ ” notation is introduced because it is more compact, especially for problems with three or more layers. Using the example above, we would write  $b^1 = a^{1|2|3}$  and  $\mathcal{B}^{(1)} = \mathcal{A}^{(1)} \times \mathcal{A}^{(2)} \times \mathcal{A}^{(3)}$ .

An important dimension of coupling is determining the physics of what can be coupled with what. For example, you need one pilot (or crew) per aircraft, but you can put two pilots on an aircraft, with one flying the aircraft and the other repositioning (flying free). We may need three locomotives to pull a train, but put five locomotives on the train because we need to reposition power (or put only two locomotives if that is enough to move the train and we are short on power).

These issues are captured by parameters of the resources and physical laws represented in the system dynamics (see below).

### 4.1.3 A note on state spaces

It is useful to discuss the issue of state spaces, since in our problem class, there are different ways of measuring the number of states. We can say that  $a \in \mathcal{A}$  is the state of a particular resource,  $R_{at}$  is the number of resources in this state at time  $t$ , and  $R_t$  is the state of our *system* (actually, only a portion of our full system state). The mathematics of control theory and dynamic programming always refer to the state of the system. In our problems, we need to capture the state of the system, but we also find it useful to represent the state of a resource. For example, it is not unusual to have someone note that a system has a large state space. We are often more interested in the size of the *resource state space*  $\mathcal{A}$  than we are in the size of the *system state space*. Thus, if  $N = \sum_{a \in \mathcal{A}} R_{at}$  is the number of resources in the system at time  $t$ , then  $|\mathcal{A}|$  is the size of the attribute space, while the state space  $S$  is the set of all possible instances of the vector  $R_t$ . If any resource can be in any element of  $\mathcal{A}$ , then the size of our state space can be shown to be given by:

$$|S| = \binom{N + |\mathcal{A}| - 1}{|\mathcal{A}| - 1}$$

If we have an instance of a single salesman spread among 50 cities, then both our resource and our system have 50 states. If we have 100 salesman that we can distribute among 50 cities, then each resource may have 50 possible states, but our system has approximately  $6.7 \times 10^{39}$  states. This is not even a big problem. In Powell & Shapiro (1996), a problem instance is given with  $|\mathcal{A}| = 1,188,000$ , and  $N = 6000$ . Thus, we wish to avoid formulations that require working explicitly with the full *system* state variable.

Even our resource state space can become extremely large in the presence of multiple resource layers when we want to consider both primitive and composite resources. In practice, modelers often avoid the complexity of composite states by handling these states in a simple way. For example, an idle taxicab usually belongs to a fairly small set of states (perhaps the set of discrete locations). When a cab picks up a customer, he will move through a sequence of intermediate steps which include picking up the customer, and driving through a series of nodes on the network before arriving at the destination. This sequence of steps is typically handled in a simple way.

Let  $\mathcal{A}$  be the set of all possible states, including all possible composite states made up of coupled

resources. We can typically divide this state space into:

$\mathcal{A}^{(c,p)}$  = Set of *primary* states (in class  $c$ ).

$\mathcal{A}^{(c,s)}$  = Set of *secondary* states (in class  $c$ ).

where  $\mathcal{A}^{(c)} = \mathcal{A}^{(c,p)} \cup \mathcal{A}^{(c,s)}$ . We only store the vector  $R_{at}$  for  $a \in \mathcal{A}^{(c,p)}$ . Thus, we have a real incentive to keep  $|\mathcal{A}^{(c,p)}|$  as small as possible.

The choice of what is a primary versus a secondary state is up to modeling judgment. For example,  $\mathcal{A}^{(c,p)}$  might include an empty vehicle sitting idle at any location, a driver at rest at his home domicile, or a machine that just completed a job. As a rule, the simulation of a resource should proceed from one primary state to another, with secondary states serving more as intermediate accounting points. Typically, it is not possible to model the arrival of new information which requires a change in decision in between secondary states.

The distinction between primary and secondary states is a powerful modeling tool. For example, a pilot may be sitting at home, or may be coupled with an aircraft, navigator and stewarding crew half way through a trip. This is a very complex state, and if we had to enumerate them all, the problem would normally be far too large. However, we can use specialized algorithms to track a resource from one primary state to another within the setting of a subproblem. In this way, secondary states can be generated and destroyed on the fly. It is through the primary states that we communicate the impact of decisions in one subproblem on another, so primary states that have been visited normally have to be retained.

## 4.2 Aggregation

In most cases, we are not interested in all the fields, and therefore we may aggregate the data, which can be represented using an aggregation function. Since in some cases it is useful to work with multiple levels of aggregation, we define:

$$G^n : \mathcal{A} \rightarrow \mathcal{A}^{G^n}$$

where  $G^n$  is the  $n^{th}$  level aggregation of the attribute space  $\mathcal{A}$ . While it may be that  $\mathcal{A}^{G^n} \subseteq \mathcal{A}^{G^{n+1}}$ , this will not always be true (for example, a driver may have a location which is a city, and we may

aggregate this into whether he is at home or not, outcomes that are not instances of cities). Since multiple resources may aggregate into the same element, we define:

$$\begin{aligned}
 R_{a,t} &= \text{number of resources at time } t \text{ that are expected to have attribute } a. \\
 &= \sum_{r \in \mathcal{R}_t} 1_{\{G(a_r)=a\}} \\
 R_t &= \text{vector of elements of } R_{a,t}
 \end{aligned}$$

where  $1_{\{x\}}$  is the indicator variable which is 1 if  $x$  is true, and 0 otherwise. If we are using the  $n^{\text{th}}$  level of aggregation, we will refer to the vector  $R^{G^n}$ . Normally, we would have preferred to use the lowercase vector  $r_t$  to indicate the vector of resources, but this creates confusion with the elements  $r \in \mathcal{R}$ .

## 5 Processes

The process component of a DRiP includes the elements that describe how a system evolves over time, excluding the control of the process which is covered under the third dimension. Our discussion of processes covers three subdimensions: information processes, system dynamics, and the specification of constraints.

Our representation of processes requires that we first adopt a convention for labeling time. For dynamic problems, it is often convenient to index variables using a time index that is relative to the current time. We let  $t^c$  refer to *clock* time and  $t$  refer to *relative* time. Let:

$t_{current}^c$  = The *absolute clock time* corresponding to  $t = 0$ .  $t_{current}^c$  may be a piece of data, or it may be a function which, when referenced, returns the “current” time.

$t^c(t)$  = The clock time corresponding to relative time  $t$ .  
 $= t + t_{current}^c$ .

In a planning setting, the clock  $t_{current}^c$  is a fixed number. In a real-time control setting, the clock  $t_{current}^c$  is changing (possibly continuously, or perhaps only in periodic jumps, such as once an hour).

We find it useful to distinguish between the *past*, defined by events where  $t < 0$ , and the *future* represented by  $t > 0$ . With this indexing,  $t = 0$  is called “here and now.” Since we sometimes have

to refer to different points in time in the context of a single variable, we will use the indexes  $t$ ,  $t'$  and  $t''$  where it will generally be true that  $t \leq t' \leq t''$ . We often have to refer to what we know at time  $t$  about what will (or might) happen at time  $t'$ , which might be represented using  $x_{t,t'}$ . We use the notation  $x_{t,\mathcal{T}_t}$  to refer to the entire vector  $\{x_{t,t'}\}_{t' \in \mathcal{T}_t}$  for a given set of (generally contiguous) time periods  $\mathcal{T}_t$ .

Variables of the form  $x_{t,t'}$  arise frequently in dynamic models, representing information about time  $t'$  that is known at time  $t$ . In most models, time  $t = 0$ , representing here and now, whereas  $t' \in \mathcal{T}^{ph}$  represents a point within the planning horizon in the model. When this is the case, it is more compact to use the simpler notation  $x_t$  where  $t \in \mathcal{T}^{ph}$  and all information is assumed to be relative to “here and now.” This creates a potential notational ambiguity, since we also will use  $x_t$  to represent the family of vectors  $x_t = \{x_{t,t'}, t \in \mathcal{T}_t^{ph}\}$ . We suggest that in the formulation of a model, the reader should consistently adopt either a single or double time indexing style to avoid this confusion.

Our discussion requires that we define:

$\mathcal{T}_t^{ph}$  = The set of time periods in the *planning horizon* for a decision made at time  $t$ , where all the data that may directly or indirectly impact the decision at time  $t$  must fall within the planning horizon.

$\mathcal{T}_t^{ih}$  = The set of time periods in the *implementation horizon* for a decision made at time  $t$ , reflecting the set of time periods over which decisions that have been planned will be implemented without replanning with new information.

$\mathcal{T}^{sh}$  = The set of time periods in the *simulation horizon*, representing the total number of time periods over which we will be evaluating a given policy.

The definitions of these horizons imply that  $\mathcal{T}_t^{ih} \subseteq \mathcal{T}_t^{ph} \subseteq \mathcal{T}^{sh}$ . The planning horizon captures the time periods where data of any form may impact a decision at time  $t$ . The implementation horizon captures the set of time periods where *decisions* planned within those time periods may have a direct impact on a decision at time  $t$ .

Our terminology is consistent with the literature on planning horizons. We have avoided explicit reference to forecast horizons and decision horizons, which have formal definitions (see, for example, Bhaskaran & Sethi (1987), Bes & Sethi (1988), Aronson & Chen (1989), Sethi & Bhaskaran (1985)).

White (1996) discusses “decision rolls” which involve the implementation of a one-period decision (based on a  $T$  period horizon) and “horizon rolls” which involve the implementation of all decisions made within the  $T$  period planning horizon.

We have replaced the concept of a “decision roll” with an “implementation horizon.” It is customary to assume that the set of events  $|\Omega_{\mathcal{T}^{ih}}| = 1$ , which is to say that there is only one set of possible outcomes within the implementation horizon. We do not require this. It is possible to have  $|\Omega_{\mathcal{T}^{ih}}| > 1$ , but this means that we would have to determine a decision  $x_{a,d}(\omega)$  for  $\omega \in \Omega_{\mathcal{T}^{ih}}$  (since we do not allow replanning within the implementation horizon, we would have to compute a decision for each possible outcome). This, in fact, is exactly the approach used by the stochastic programming community.

## 5.1 Information processes

An explicit dimension of our representation of the system is the arrival of information to the system.

We consider two classes of information:

- 1) Exogenous information processes (new information arriving to the system).
- 2) Endogenous information processes (internal decisions).

Although the process of making decisions is covered under the dimension “Controls,” we need to capture the effect of a decision on our process. We found that there are parallels between exogenous information and decisions that seem to be overlooked in the classical modeling literature.

### 5.1.1 Exogenous information processes

We can divide exogenous information processes into two broad classes:

- 1) Information generated by an *actual* exogenous process.
- 2) Information generated by a *forecasted* exogenous process.

From the perspective of building a model, we deal only with forecasted processes. These are generally much simpler than actual processes. For example, we may wish to model weather delays to aircraft, but not mechanical failures or the fact that a crew may not show up for work.

The only information that we can model is information that updates an element within our system. Recall that  $\mathcal{C}^I$  represents our set of information classes in our system and that  $\mathcal{E}_t^{(c)}$  is the set of all information elements in class  $c$  at time  $t$ . To model the process of information arriving to the system, let:

$\hat{\mathcal{E}}_t^{(c)}$  = The set of new information elements arriving at time  $t$  for class  $c$ . If  $e \in \hat{\mathcal{E}}_t$  but  $e \notin \mathcal{E}_t$  then this represents a new information element entering the system.

$\kappa_{et}$  = An individual data packet arriving at time  $t$  to update information element  $e \in \hat{\mathcal{E}}_t$ .

$\kappa_t$  = The collection of all information arriving at time  $t$ .

$$= \bigcup_{e \in \hat{\mathcal{E}}_t} \kappa_{et}$$

An information packet contains two types of information. First, we can think of an information packet for element  $e$  as consisting of a revised set of elements of the vector  $a_e$  (in general, a packet will consist of an update of only a subset of these elements). The second piece of information is when this data will become *effective*, which we define as:

$t_{et}^{eff}$  = The time at which the information in  $\kappa_{et}$  becomes *effective*, which is to say the first time that it can be used in the system dynamics (see below) to describe the evolution of the system. We assume that  $t_{et}^{eff} \geq t$ .

(A more general representation would provide for a  $t^{eff}$  for each element in  $a_e$ ; our representation will do for now.) Following our convention earlier (see section 4), we let  $a_{time_e} = t^{eff}$  store the time at which the data in  $a$  becomes effective.

As each new piece of information arrives, we assume that there is a process for updating the database in some way. This can be generally represented using:

$$K' \leftarrow U^K(K, \kappa_e) \tag{1}$$

or:

$$K_{t+1} = U^K(K_t, \kappa_t) \tag{2}$$

The update function is easiest to understand using the version in (1). If the element  $e \in \mathcal{E}$ , then typically the update function is taking the new elements in  $\kappa_e$  and modifying the appropriate

elements in  $a_e$ . If  $e$  is a new element, then the update function has to augment the set of elements  $\mathcal{E}$  and then add the new element. Similarly, we assume that there is a protocol for indicating when an information element  $e$  needs to be dropped from the set  $\mathcal{E}$ .

In general, the sequence  $\kappa_t$  will be random, so, using standard convention, we define:

$\Omega =$  The set of possible outcomes of data packets  $\kappa_t, t \in \mathcal{T}$ , for points in time  $t \geq 0$ .

**Remark:** At this point, it is useful to contrast our representation with classical stochastic processes. Let  $h_t = \{\kappa_0, \kappa_1, \dots, \kappa_{t-1}\}$  be the history of the (forecasted) process, giving all the events prior to time  $t$  (but after time 0). Assume that  $h_t \in \mathcal{H}_t$  and let  $\mathcal{F}_t = \sigma(H_t)$  be the  $\sigma$ -algebra on the set  $\mathcal{H}_t$ . As a rule,  $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$  which means that events up to time  $t$  are not forgotten by time  $t + 1$  (since these events remain in the information set). In this case,  $\mathcal{F}_t$  is a *filtration* (hence the notation  $\mathcal{F}_t$ ). In real systems, this standard assumption is often violated. The update process  $U^K(K_t, \kappa_t)$  typically *does* in fact “forget” past data. It is not unusual for corporate databases to retain detailed records of only a limited subset of past history. For this reason, our notation  $K_t$  more accurately depicts the real information flow.

The sample space  $\Omega$  must be defined over some horizon. When necessary, we use notation such as  $\Omega_\infty$  to denote an infinite horizon, or  $\Omega_{\mathcal{T}}$  to denote a sample space over a set of time periods  $\mathcal{T}$ . An element  $\omega \in \Omega_{\mathcal{T}}$ , then, represents a particular outcome of the sequence  $\{\kappa_t, t \in \mathcal{T}\}$ .

The vector  $K_t$  captures the data about all the objects in the system. The subset of information about resources is, for the most part, contained in the vector  $R_t$ . In some cases, however, it is notationally convenient to have a global set  $K_t$  to describe all the data for the system. The data in  $K_t$  that is not captured by  $R_t$  primarily covers parameters used in the modify function (defined below).

We now wish to distinguish between actual and forecasted processes. The sequence of data packets  $\kappa_t, t < 0$  represents information that has actually happened, whereas  $\kappa_t, t > 0$  represents information that might happen in the future. While these two sets should be defined with respect to the same probability space, in practice this will never be true. (We assume that “actual” events are drawn from a probability space  $(\Omega^a, \mathcal{F}^a, \mathcal{P}^a)$ .) Similarly, simulated events are defined with respect to a probability space  $(\Omega, \mathcal{F}, \mathcal{P})$ . Also, we will always need to estimate the probability measure  $\mathcal{P}$ , whereas we will generally not need (nor will we be able) to estimate  $\mathcal{P}^a$ .

An important difference between actual and forecasted events is that we have no control over what actually happens, but we can make simplifying assumptions with regard to what we are forecasting will happen in the future. For example, while we may receive actual updates on both resources and parameters, we generally only forecast updates to resources (customer orders being called in, changes in driver ETA's, and so on). Other data may also be changing (for example, the price of fuel) but we may not want to model these changes in the future.

The challenge of forecasting is to estimate the set of outcomes  $\Omega$  and the probability measure  $\mathcal{P}$ . For simplicity, assume that outcomes are discrete, so that for  $\omega \in \Omega$ , the probability  $p(\omega)$  is well defined. There are four strategies that are generally used to create a probability law for the forecasted sample space  $\Omega$ :

1. We ignore future events ( $\Omega$  is empty).
2. We use a point estimate of the future, which is assumed to occur with probability one (in this case, we assume that  $|\Omega| = 1$ ).
3. We use past history to estimate the parameters of an assumed probability distribution (e.g. normal, Poisson). In this case  $p(\omega)$  is derived from the assumed probability distribution, from which we can calculate expectations or randomly sample outcomes.
4. We use a set of samples drawn from past history to create a (finite) sample space of future outcomes (typically,  $p(\omega) = 1/|\Omega|$ ).

Myopic models use the first option, deterministic models use the second option, and most stochastic models use the third. The stochastic programming community (and especially the modeling of financial problems) uses the fourth option, primarily because there is not an accurate mathematical model of the underlying physical processes.

Both deterministic models and stochastic models require the ability to estimate parameters from which a probability distribution may be constructed. To formalize this process, define:

$\rho_{t,t'}$  = Vector of parameters at time  $t$  to be used to forecast a realization of a random outcome at time  $t'$ .

$\rho_{e,t,t'}$  = Subvector of  $\rho_{t,t'}$  used to estimate at time  $t$  the outcome of element  $e \in \mathcal{E}$  at time  $t'$ .

$$\rho_t = \{\rho_{t,t'}, t \in \mathcal{T}_t^{ph}\}$$

The subvector  $\rho_{e,t,t'}$  might be the mean and variance of a travel time or a fuel price. We may wish to represent this outcome explicitly as a normal distribution, or simply use a point estimate (such as the 80<sup>th</sup> percentile estimate of the travel times).

We have written the estimation of our parameters as  $\rho_{t,t'}$ , but in practice, we typically only have estimates as of time  $t = 0$ . We could, in theory, update  $\rho$  as simulated events occur (that is,  $\kappa_t$  for  $t > 0$ ), but this is not normal modeling practice.

We represent the process of estimating the forecasting parameters using another class of update function, which we denote  $U^f$ . The update process would be written:

$$\rho_{t+1} = U^f(\rho_t, \kappa_t)$$

We close this discussion by observing that there are two views of the future that need to be represented. We have represented  $\Omega$  as the set of outcomes that we are forecasting may happen in the future. This set may include the possibility of a number of extreme outcomes that we do not necessarily want to consider when developing a plan. For this reason, we define:

$$\Omega^p = \text{The set of outcomes that we wish to } \textit{plan} \text{ on during the horizon } \mathcal{T}^{ph}.$$

We give the set  $\Omega^p$  a special name because it is the set of outcomes that a company has planned on when developing a strategy. This set may contain a single outcome, or a set of “anticipated outcomes.” This representation allows us to capture the behavior of companies to plan for certain contingencies, but not all of them. The set  $\Omega^p$  may be the same as  $\Omega$ , but not necessarily. If they are different, then we are planning around a set of anticipated outcomes (or a single outcome), but we may still want to evaluate a solution in terms of a larger set of outcomes.

### 5.1.2 Endogenous information processes

In addition to exogenous information, our system evolves as a result of endogenous information, in the form of a sequence of decisions. A thorough discussion of decisions and other methods of control is given in the third dimension of our paradigm. Here, we are not so concerned with the issue of how decisions are made, but rather with the types of decisions and their effect on our system.

Let:

$\mathcal{C}^D$  = Set of decision classes.

$\mathcal{D}^{(c)}$  = Set of specific decisions in class  $c \in \mathcal{C}^D$ .

$$\mathcal{D} = \bigcup_{c \in \mathcal{C}^D} \mathcal{D}^{(c)}$$

$d_{t,t'}$  = A vector of decisions made in time  $t$  to be implemented at time  $t'$ .

$$d_t = \{d_{t,t'}, t' \in \mathcal{T}_t^{ph}\}$$

We call  $d_t = \{d_{t,t'}, t' \in \mathcal{T}_t^{ph}\}$  a *plan*, just as we call  $\omega = \{\kappa_t, t \in \mathcal{T}^{ph}\}$  a *forecast*. Viewed in this way, we can think of a *plan* as a *forecast* of future decisions. For example, an airline might have the control classes:

$$\mathcal{C}^D = \{\text{move aircraft, purchase/retire aircraft, accept a reservation, set ticket prices, cancel a flight}\}$$

The class “move aircraft” might include decisions:

$$\mathcal{D}^{\text{move aircraft}} = \{\text{Chicago, Detroit, Newark, ...}\}$$

It is often convenient in practice to define:

$\mathcal{D}^{(c,a)}$  = The set of decisions in class  $c$  that can be applied to resources with attribute  $a$ .

For accounting purposes, we need to *count* the number of times a particular transformation is made.

For this reason, we define the *counting variables*:

$$x_{a,d,t} = \text{Number of times decision } d \text{ was executed at time } t \text{ on a resource with attribute } a, \text{ where } x \text{ need not be integer.}$$

We sometimes use the notation  $x_{a,d,t}(\omega)$  when we want to express the dependence of the decisions  $x_t$  on a specific set of outcomes.

**Remark:** It is useful to briefly contrast our notation with classical modeling techniques. First, we note that it is more common to represent flow, say from node  $i$  to node  $j$ , as  $x_{ij}$ . This would be comparable to writing our counting variable as  $x_{a,a'}$  (letting the time dimension be implicit in

$a$  and  $a'$ ). In simple problems, the  $x_{ij}$  representation works because we think of the *decision* as going from  $i$  to  $j$ . In our more complex problems, the *decision* might be to move from location  $i$  to location  $j$ , but this decision carries implications on the evolution of other dimensions of the attribute vector. Thus, we do not really *decide* to go from state  $a$  to state  $a'$ , but rather we are in state  $a$ , and decide to implement an action  $d \in \mathcal{D}$ . The outcome  $a'$  may be uncertain, so it is inappropriate to write  $x_{a,a'}$  unless we want to condition on the outcome.

Aside from the source of the information, the biggest difference between an exogenous information event  $\kappa_t$  and an endogenous decision  $d_t$  is the means by which the information changes the knowledge base. In the case of exogenous information, the update is handled through the update function  $U^K$ . Typically, we assume that a data packet  $\kappa_{et}$  is directly revising elements of an attribute vector  $a_e$ . In the case of a decision, the impact on the system is handled through the system dynamics, described later.

It is useful to consider the similarity between an endogenous decision, which changes a resource, and an exogenous data event that changes a resource. Since these both have the same effect, it is useful to define:

- $d_{et}^0$  = Changes to a resource in the data packet  $\kappa_{et}$ ,  $e \in \mathcal{E}^{(c)}$ ,  $c \in \mathcal{C}^R$ .
- $x_{a_e, d_{et}^0, t}^0$  = Number of times an exogenous change  $d_{et}^0$  acted on a resource  $e$  with attribute  $a_e$ .
- $x_t^0$  = Vector of all exogenous changes to resources in time period  $t$ .

Our choice of notation  $d_{et}^0$  is motivated by the similarity between endogenous controls, represented by  $d_t$  which have the effect of changing  $R_t$ , and exogenous changes that accomplish the same function. (We note that the notation  $d$  and  $d^0$ , motivated by the mnemonic “decision”, also has the convenient mnemonic “delta” capturing a change to the system.) We similarly define  $d_t^0$  as a set of exogenous changes occurring in time period  $t$ .

The concept of a “planned” set of (possibly random) outcomes is widely used in practice, although in these cases the plans are generally deterministic. Often, along with the set  $\Omega^p$  we may also define:

- $x_{t, \mathcal{T}_t}^p$  = A plan, made at time  $t$ , in the form of a set of decisions over a set of time periods  $\mathcal{T}_t$ .

A plan  $x^p$  is often a set of decisions that has been developed at an earlier time and communicated to individuals that are exogenous to our system. For this reason, deviations from plan come with a cost. If  $|\Omega^p| > 1$ , we might specify a plan for each outcome, that is,  $x^p(\omega)$  for  $\omega \in \Omega^p$ .

### 5.1.3 A note on model updating

It is important to distinguish between the updating of information due to actual exogenous processes (which by definition occur in the past) from the updating of information due to simulated processes in the future. The issue is most significant in the context of *forward* algorithms that start at time  $t = 0$  and progress forward in time, using a Monte Carlo sample. Such procedures are used in both simulation software, and are becoming increasingly popular in problems that use dynamic programming (see Sutton & Barto (1998) and Bertsekas & Tsitsiklis (1996)) and Bender's decomposition in stochastic programming (Higle & Sen (1991), Infanger & Morton (1996), Morton (1996) and Pereira & Pinto (1991)).

When we are working in the future, it is common to simulate a new event that we may reference as  $\kappa_t(\omega)$  (the dependence on  $\omega$  indicates that this is a random sample), and then update the knowledge base:

$$K_{t+1}(\omega) = U^K(K_t(\omega), \kappa_t(\omega))$$

but we generally would not update the forecast parameters  $\rho$ , since these would have been estimated using purely historical data.

## 5.2 System dynamics

There are two ways that a resource can be modified: exogenous controls, captured by  $\kappa_t$  (or  $d_t^0$ ), and endogenous controls, represented by a sequence of decisions  $d_t$ . While any piece of exogenous information can change our system, we assume that only changes in the status of resources, represented by  $d_t^0$ , can produce a direct change. Changes in information classes that are not resources may trigger a decision  $d_t$ , and therefore produce a change. As a result, exogenous changes in resources  $d_t^0$  and decisions  $d_t$  play very similar roles, and it is useful to think of these as, respectively, exogenous and endogenous controls. Throughout the remainder of this section, we refer to decisions  $d_t$ ; in each case, the reader may substitute exogenous controls  $d_t^0$ .

At the heart of our DRiP is a function, which we call the modify function, which captures the changes induced by a control. This function can be represented as the mapping:

$$M_t(a, d_t, K_t) \rightarrow (a', c, \tau) \tag{3}$$

where:

$a$  = The attribute vector of the (possibly composite) resource being modified.

$d_t$  = An endogenous (or exogenous) control.

$K_t$  = The knowledge base of our system at time  $t$ .

$a'$  = The attributes of the modified resource.

$c$  = The (possibly vector of) costs (or contributions) resulting from the action.

$\tau$  = The *transfer* time, giving the time required to complete the action.

(Note the convenient mnemonic that the right hand side of (3) spells “act.”) This represents a modification of an attribute at time  $t$  based on a decision to be implemented at time  $t$ . Often, we are planning a decision  $d_{t,t'}$  at time  $t$  to be implemented at time  $t'$ . In this case, we would write the mapping:

$$M_t(a, d_{t,t'}, K_{t,t'}) \rightarrow (a', c, \tau) \tag{4}$$

The knowledge base  $K_t$  depends, of course, on the flow of exogenous information  $\{\kappa_t, t \in \mathcal{T}\}$ . An instance of this sequence of information is captured by an element  $\omega \in \Omega$ , meaning that we can equivalently write (3) as:

$$M_t(a, d_t, \omega) \rightarrow (a', c, \tau) \tag{5}$$

This way of writing the system dynamics is popular in certain parts of the modeling literature. It is both mathematically accurate and notationally more compact, but it does not express as precisely the information content of a decision as the forms expressed in (3) or (4). The applied probability community will use the representation in (5), where it is “understood” that the modify function uses only the history of  $\omega$  up to time  $t$ . It is somewhat more precise to write  $M_t(a, d, h_t)$  where  $h_t$  is the history of the process (some communities use  $\omega^t$  to capture the history up to time  $t$ , but

this runs against our notational style), or  $M_t(a, d, \mathcal{H}_t)$ , where  $\mathcal{H}_t$  is the  $\sigma$ -algebra generated by  $h_t$  (many people use  $\mathcal{F}_t$  instead of  $\mathcal{H}_t$ ). As we observed earlier, this carries the implicit assumption of a memoryless process which is independent of the updating mechanism  $U^K$ .

We use both systems in our discussion; we prefer the representation in (3) or (4) when the additional precision is needed, and we use (5) when we simply want to capture the dependence of a transformation on exogenous information.

If a resource is a composite resource, then this fact is reflected in the elements of the vector  $a$  which indicates the presence of coupled resources. The modify mapping must be coded to check for the presence of coupled resources and update the attributes of these resources as well.

We generally need to know how a decision actually impacts the system. Let:

$$\begin{aligned} \delta_{a',t,t'}(a, d, \omega) &= \text{Change in the system at time } t' \text{ given a decision executed at time } t. \\ &= \begin{cases} 1 & \text{if } M_t(a, d_t, \omega) = a', t + \tau = t' \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Note that we represent the triplet  $(a, d, \omega)$  as an argument for  $\delta$ , because we think of  $\delta$  as a function with argument  $(a, d, \omega)$ .

We may use the  $\delta$  function to represent the impact of both endogenous decisions and exogenous changes on our resource vector. Thus,  $\delta_{a',t}(a_e, d_{et}^0, \omega)$  would be the impact of an endogenous change  $d_{et}^0$ , acting on element  $e$ , on resources with attribute vector  $a'$ .

Using this notation, we can express the evolution of the system using:

$$\begin{aligned} R_{a,t,t'}(\omega) &= R_{a,t-1,t'}(\omega) + \sum_{e \in \hat{\mathcal{E}}_t} \delta_{a',t,t'}(a'_e, d_{et}^0, \omega) x_{a'_e, d_{et}^0, t}^0 \\ &\quad + \sum_{d \in \mathcal{D}} \sum_{a' \in \mathcal{A}} \delta_{a',t,t'}(a', d_{t,t'}, \omega) x_{a', d, t} \quad a \in \mathcal{A}, \quad t' \geq t \end{aligned} \tag{6}$$

Equation (6) contains three elements on the right hand side: 1) what we knew about the status of resources for time  $t'$  at time  $t - 1$ , 2) the exogenous changes to resources (for time  $t'$ ) that arrived during time  $t$ , and 3) the endogenous changes to resources for time  $t'$  that were made during period  $t$ .

### 5.3 Constraints

In section 2.2, we divided constraints between conservation of mass (or conservation of flow) and constraints on the rate of process transformation. Conservation of flow at each point in time  $t$  is easily written as:

$$\sum_{d \in \mathcal{D}} x_{a,d,t} = R_{at} \quad \forall a \in \mathcal{A} \quad (7)$$

We listed three types of constraints that restrict the rate of process transformation. All three of these can be modeled in the same way. Using the counting variable  $x_{a,d,t}$  we can express constraints on the rate of process transformation. We might, for example, require that:

$$x_{a,d,t} \leq u_{a,d,t}$$

where  $u_{a,d,t}(\omega)$  may represent a physical constraint (the capacity of a truck or terminal) or a market demand (the transition  $a \rightarrow a'$  may represent satisfying a demand, where  $u_{a,d,t}(\omega)$  is the amount of the demand). As we point out later, we can only represent a market demand as a constraint if unsatisfied demands are lost. Note that we allow the bound to be a function of both endogenous and exogenous controls.

In some cases, a constraint is specified at an aggregate level. Given subsets  $\mathcal{A}'$ ,  $\mathcal{T}'$  and  $\mathcal{D}'$ , we may write:

$$x_{\mathcal{A}', \mathcal{D}', \mathcal{T}'}(\omega) = \sum_{a \in \mathcal{A}'} \sum_{d \in \mathcal{D}'} \sum_{t \in \mathcal{T}'} x_{a,d,t}(\omega)$$

then we may require that:

$$x_{\mathcal{A}', \mathcal{D}', \mathcal{T}'}(\omega) \leq u_{\mathcal{A}', \mathcal{D}', \mathcal{T}'}(\omega)$$

It is extremely valuable, perhaps even essential, in a stochastic control problem that a feasible solution always exist. For this purpose we make the following assumption:

**Assumption 1:** If  $a \in \mathcal{A}^{(p)}$  (a primary state), then  $d^\phi \in \mathcal{D}^{(a)}$ .

This allows us to assume that the *hold* option always exists from a primary state. Then, we impose:

**Assumption 2:** Given any finite upper bound  $u_{\mathcal{A}', \mathcal{D}', \mathcal{T}'}(\omega)$  we require that  $d^\phi \notin \mathcal{D}'$ .

Thus, we cannot limit our ability to hold (in a primary state). In our work, we have never found it necessary to use a hard constraint on our ability to hold resources in a primary state. By requiring that the hold decision be unconstrained (although it may face cost and time penalties) we guarantee that there will *always* be a feasible solution to any DRiP.

The reader will notice that we have pointedly left out constraints such as time window constraints. It is our belief that the phrase “time window constraint” is a phrase based on the vocabulary of linear programming. Time windows arise as a result of a reward structure that varies over time. In static problems, one way of capturing this nonlinear structure is through a constraint. Our belief is that time windows are not true physical constraints (that is, we do not violate any physical laws if a time window is violated), and therefore should be modeled through the reward structure, and not as a constraint.

## 6 Controls

The basic types of decisions has already been introduced in section 5.1.2 under the heading of endogenous information processes. Now we turn to the issue of describing *how* decisions are made. This discussion is divided into four components: the control structure, the information set, the design of the decision function, and measurement and evaluation.

### 6.1 Control structure

Large, complex systems invariably exhibit a multi-agent control structure. We propose that any set of decisions  $\mathcal{D} = \bigcup_{c \in \mathcal{C}^D} \mathcal{D}^{(c)}$  can be broken along three dimensions:

- 1) A subset of the attribute space. This subset may be based on geography, the type of resource being managed, and the status of the resource.
- 2) A subset of the control classes which are the responsibility of an agent.
- 3) A subset of time periods within the planning horizon.

We define the decomposition of the problem using:

$\mathcal{Q}$  = Set of subproblems that encompass the problem.

$\mathcal{A}_q$  = Subset of the attribute space for subproblem  $q$ , where  $\bigcup_{q \in \mathcal{Q}} \mathcal{A}_q = \mathcal{A}$  and  $\mathcal{A}_{q_1} \cap \mathcal{A}_{q_2} = \phi$  when  $q_1 \neq q_2$ .

$\mathcal{C}_q^D$  = Set of control classes associated with subproblem  $q$ .

$\mathcal{T}_q^{ih}$  = The *implementation horizon* for subproblem  $q$ . This is the set of time periods during which subproblem  $q$  controls the decisions.

The information elements within our control would be given by:

$$\mathcal{E}_q = \{e | a_e \in \mathcal{A}_q\}$$

If we let  $d_t$  be a decision to be implemented at time  $t$ , then we can define the set of decisions in subproblem  $q$ :

$$\begin{aligned} \mathcal{D}_q &= \text{Subset of decisions in subproblem } q. \\ &= \{d_{t'} \in \mathcal{D}^{(a,c)}, c \in \mathcal{C}_q^D, a \in \mathcal{A}_q, t' \in \mathcal{T}_q^{ih}\} \end{aligned}$$

Recall that we may make a decision at time  $t$  to be implemented at time  $t'$ . The element  $q$  implies the implementation time, but not the time the decision is made. For this reason, we let:

$$\mathcal{D}_{t,q} = \text{Set of decisions that are planned at time } t \text{ in the set } \mathcal{D}_q \text{ to be implemented (by definition of } q) \text{ within the interval } \mathcal{T}_q^{ih}.$$

When we divide a problem into subproblems, we need to establish a few concepts establishing the relationship between subproblems. First define:

**Definition 6.1** *We say that two decisions in the same subproblem are **tightly coupled**. Two decisions that impact each other (perhaps indirectly) which are not in the same subproblem are **loosely coupled**.*

We especially need notation that describes what subproblems are affected by decisions in other subproblems. For this we define:

**Definition 6.2** *The forward reachable set is given by:*

$$\vec{\mathcal{M}}_q = \{q' \in \mathcal{Q} \mid \exists a \in \mathcal{A}_q, d \in \mathcal{D}_q \text{ with } M(a, d, \cdot) \mapsto (a', \cdot, \cdot), a' \in \mathcal{A}_{q'}\}.$$

Thus, the forward reachable set is the set of subproblems  $q'$  which can be reached directly from  $q$ . Conversely we may define:

**Definition 6.3** *The backward reachable set is given by:*

$$\overleftarrow{\mathcal{M}}_q = \{q' \in \mathcal{Q} \mid q \in \vec{\mathcal{M}}_{q'}\}$$

An important special case of a control structure is where  $\mathcal{Q} = \mathcal{T}$ , which is to say that subproblem  $q$  is equivalent to time  $t$ . Frequently, we will address a problem from the perspective of a single subproblem (in the sense that we are looking from the perspective of a single agent), but we still have to model the evolution of information over time. In this special case,  $\vec{\mathcal{M}}_q = \vec{\mathcal{M}}_t = t + 1$ .

## 6.2 The information set

Up to now, we have used the word “information” in a relatively general way. As preparation for building a decision function, we have to specify the information that is available to us when we make a decision.

In our view, there are two broad classes of information: exogenous, which is given to us before we start modeling a physical system, and endogenous, representing information that we generate internal to our system. Exogenous information, as was discussed in section 2.1, can be divided between information giving the state of the system, which we called  $K_0$  (what we know at time  $t = 0$ ), and (optionally) a plan, represented by  $x_0^p$ , providing an indication of decisions to be made in the future. Endogenous information is, strictly speaking, a type of forecast: projections of events that we believe will happen in the future, but which we have to generate endogenously (because they have not happened yet). We may forecast updates to  $K_0$  (captured by the set  $\Omega$ ) or new decisions that will be made (determining these decisions is, ultimately, the goal of the modeling process).

We find it useful to think of  $K_0$  as “what we know now,”  $\Omega_{\mathcal{T}^{ph}}$  as the forecast of future exogenous events, and the plan  $x_{\mathcal{T}^{ph}}^p$  as the “forecast” of future decisions (keep in mind that it is common in

many models to ignore the presence of  $x^p$ , even if it exists). In this sense, it seems natural to group  $\Omega_{\mathcal{T}^{ph}}$  and  $x_{\mathcal{T}^{ph}}^p$  together as the forecast of future exogenous and endogenous information processes.

At this point, then, we have identified two classes of information that we may (or may not) use in our decision function at time  $t$ : the state of the system  $K_t$ , and forecasts of exogenous and endogenous processes,  $\Omega_{\mathcal{T}_t^{ph}}$  and  $x_{\mathcal{T}_t^{ph}}^p$ . There is, however, a third class of information that can be found in several types of algorithms which captures the *value* of decisions made now on the future or, more generally, on other subproblems. Let:

$V_{q'}(R_{q'})$  = The value of resource vector  $R_{q'}$  to subproblem  $q'$ .

$$\begin{aligned} \mathcal{V}_{\mathcal{M}_q}^{\rightarrow} &= \{V_{q'}(R_{q'})\}_{q' \in \mathcal{M}_q} \\ &= \text{The family of value functions impacted by decisions made in subproblem } q. \end{aligned}$$

The third class of information,  $\mathcal{V}_{\mathcal{M}_q}^{\rightarrow}$ , is widely used in dynamic programming (in the form of a value function) and stochastic programming (via recourse functions or approximations of the recourse function such as Bender's cuts). If  $\Omega_q$  is a forecast of future (primal) events, and  $x^p$  is a forecast of future decisions, then  $\mathcal{V}_{\mathcal{M}_q}^{\rightarrow}$  can be thought of as a forecast of dual variables. In dynamic programming, a value function can be computed using classical backward techniques (e.g. Puterman (1994)) or using adaptive methods based on Monte Carlo sampling within forward dynamic programming (Bertsekas & Tsitsiklis (1996), Sutton & Barto (1998)). In stochastic programming, it has been popular to "approximate" the future using Bender's cuts (Van Slyke & Wets (1969), Birge (1985), Higle & Sen (1991), Chen & Powell (to appear)) which also has the effect of approximating (in a formal way) the impact of a decision now on the future. Other authors have developed specialized approximations for resource allocation problems. Powell & Carvalho (1998) and Powell & Shapiro (1996) use linear approximations, while Cheung & Powell (1996) and Godfrey & Powell (1999) suggest separable nonlinear approximations.

We have, then, three classes of information that may be used in a decision function. We let  $\mathcal{I}$  represent all the different classes of information sets, and define three subclasses:

$$\begin{aligned} \mathcal{I}^M &= K_t \\ &= \text{The class of myopic policies.} \end{aligned}$$

$$\mathcal{I}^{PF} = \{K_t, (\Omega_{\mathcal{T}_t^{ph}}, x_{\mathcal{T}_t^{ph}}^p)\}$$

= The class of policies using a primal forecast (which may be divided with or without the plan  $x^p$ ).

$$\mathcal{I}^{VF} = \{K_t, (\Omega_{\mathcal{T}_t^{ph}}, x_{\mathcal{T}_t^{ph}}^p), \mathcal{V}_{\vec{\mathcal{M}}_q}\}$$

= The class of policies using a value forecast. Typically, this class of policies includes forecasts and plans over time periods in the set  $\mathcal{T}_q$  (which might be a single time period).

In a multiagent setting, we would define an information set  $\mathcal{I}_q$ , representing the information available to subproblem  $q$ . In this context, the value functions  $\mathcal{V}_{\vec{\mathcal{M}}_q}$  are already quite well defined. We would typically define the set  $\Omega_q$  to capture the forecasted updates to the information elements  $\mathcal{E}_q$ , and the plan  $x_q^p$  would include only decisions (or aggregations of decisions) in the set  $\mathcal{D}_q$ . The design of the set  $K_q$  is, however, perhaps somewhat less obvious. We could of course constrain  $K_q$  to the elements in  $\mathcal{E}_q$ , but it is easy to think of operational settings where an agent  $q$  at time  $t = 0$  (here and now) might reasonably use information available in a common database about resources under someone else's control. It is up to the modeler to determine what information a decision maker may reasonably consider in a decision, and we allow multiple decision makers to have access to information about resources that are not under their control.

Having defined the information set that is available to us, we have to specify how the information is updated so that we can model the information set  $I_t$  for  $t \in \mathcal{T}^{ph}$  (that is, within our planning horizon). Already, we have defined the update function  $U^K$  that specifies how the database  $K_t$  is updated with new information. Also, we have specified the updating of the forecasting parameters using the function  $U^f$  (we assume no change in the basic functional form). As a general rule, we do not model the updating of the plan  $x^p$  as part of the forecasting process (this is typically fixed at a given point in time). This leaves only the updating of the value functions.

There are two ways of estimating  $\hat{V}_t$ : statically and dynamically. A static approximation is created once (for a given time horizon). This might be done via classical backward dynamic programming, or, in the context of multistage linear programming, by using a complete set of scenarios to create Bender's cuts (as in Van Slyke & Wets (1969) and Birge (1985)). Under such a scheme, the functions  $\hat{V}_t$  would only be updated when new information is available from history (that is, it is not updated iteratively using purely forecasted information). In such a scheme, the updating of  $\hat{V}$  is similar to the updating of the forecast parameters  $\rho$ . With respect to a specific history  $h_t, t < 0$ , then, the value functions would be static.

The second method is where the functions are updating dynamically within the algorithm. Such techniques arise in the context of forward dynamic programming (Bertsekas & Tsitsiklis (1996), Sutton & Barto (1998)) and nested Bender’s decomposition (Higle & Sen (1991), Birge (1985), Infanger & Morton (1996), Pereira & Pinto (1991)). Using these methods, we assume that each solution of the decision function  $x_t^\pi(K_t, \omega_t, \hat{V}_{t+1})$  produces a vector  $\hat{v}_t(\omega)$  that can be used to update  $\hat{V}_t$  using:

$$\hat{V}_t \leftarrow U^V(\hat{V}_t, \hat{v}_t)$$

The vector  $\hat{v}_t$  might be a vector of estimates of the value of a particular state, or it might be a vector of duals used in the generation of Bender’s cuts.

With the addition of the function  $U^V$ , we now have our full family of information updating functions:

$$\mathcal{U} = (U^K, U^f, U^V)$$

We now have the basis for updating our information set if new exogenous or endogenous information arrives to our system. This can be expressed using the general mapping:

$$\mathcal{U} : \mathcal{I} \times \mathcal{K} \mapsto \mathcal{I}$$

### 6.3 The decision function

The next step is to define *how* we make decisions within a subproblem. For this purpose, we define:

$$\begin{aligned} x_q &= \text{Vector of decisions for subproblem } q, \\ &= \{x_{dq} | d \in \mathcal{D}_q\} \end{aligned}$$

$x_q^\pi(I_q)$  = The function that determines the set of decisions  $x_q$ , where:

The information field  $\mathcal{I}_q$  defines the information that the decision function  $x_q^\pi$  has access to (we can think of it as the “IQ” of the decision function). The set  $I_q$  is a specific dataset, and is therefore a random variable. For this reason, we may write  $I_q(\omega)$  to express this dependence on the sequence of outcomes. This notation elegantly communicates the relationship that the information set is a random function of the input datastream, and the decision function depends on the information set (and is therefore also random).

The superscript  $\pi$  is derived from the term “policy” in dynamic programming. When optimizing our system, our challenge is often one of finding the best function  $x_q^\pi$  for subproblem  $q$ . So, we can think of  $\pi$  as expressing both the structure of the function (the specification) and any parameters in the specification. Let:

$\Pi =$  The set of all possible decision functions that may be specified.

The reader will probably find it useful to separate classes of policies versus instances of a policy within a class (which would require setting all the parameters within the policy).

We can divide the class  $\Pi$  into three major subclasses:

$\Pi^M =$  The class of myopic policies, where  $I_t = K_t$ .

$\Pi^{PF} =$  The class of policies using a primal forecast, where  $I_t = (K_t, \Omega_{\mathcal{T}^{ph}})$ . Note that if  $|\Omega_{\mathcal{T}^{ph}}| = 1$ , then we get a classical deterministic, rolling horizon policy which is so widely used in dynamic settings.

$\Pi^{VF} =$  The class of policies using a value forecast, where  $I_t = (K_t, \Omega_{\mathcal{T}^{ph}}, \mathcal{V}_{\mathcal{M}_t}^{\rightarrow})$ , where we use a functional approximation to capture the impact of decisions at one point in time (or one subproblem) on the rest of the system.

The specification of the information content of a decision appears to be new in the modeling literature. Deterministic models assume that each decision contains all the information available in the system; stochastic models typically assume that a decision has access to all static parameters, and the history of any stochastic processes up to time  $t$  when a decision is made. In complex systems such as freight transportation, it is common for one decision maker to simply not have access to information that are available to others simply because of the way the computer system is designed, or as a result of his own work process.

We can provide some structure to the information field  $\mathcal{I}_q$ . In a multiagent system, we can define three classes of information that might be included in  $\mathcal{I}_q$ :

$$\mathcal{I}_q = \{ \mathcal{K}_q, (\Omega_q, \mathcal{X}_q^p), \mathcal{V}_{\mathcal{M}_q}^{\rightarrow} \}$$

where:

$\mathcal{K}_q =$  The information elements that are used in subproblem  $q$ , derived directly from the sequence  $\kappa_t$  and stored in the database using the update function  $U^K$ .

$\Omega_q =$  Information events that we have forecasted for time periods  $\mathcal{T}_q$  based on parameters estimated from historical events.

$\mathcal{X}_q^p =$  The “plan” for subproblem  $q$ , representing a projected set of future decisions. In general, there may be a separate plan for each scenario  $\omega_q \in \Omega_q$ , but in practice there will typically be a single plan that we may denote  $x_q^p$ .

$\mathcal{V}_{\vec{\mathcal{M}}_q} =$  A forecast of the economic impact of decisions made in subproblem  $q$  on subproblems in the forward reachable set for  $q$ .

It is not necessary to use all three classes of information. Myopic models use  $\mathcal{I}_q = \mathcal{K}_q$ . Rolling horizon models which use a forecast of future events use  $\mathcal{I}_q = \{\mathcal{K}_q, \Omega_q\}$  (where, once again,  $|\Omega_q| = 1$  means we are using a deterministic approximation of the future, while  $|\Omega_q| > 1$  means we have a stochastic model). Dynamic programming algorithms and some stochastic programming algorithms use a value function of some sort.

## 6.4 Measurement and evaluation

Our last task is to compare one solution to another. That is, given two decision functions  $x^{\pi_1}$  and  $x^{\pi_2}$ , how do we know which is better? In this section, we discuss how to measure the performance of a function, focusing on the choice of sample space. This presentation is undertaken from two perspectives. Section 6.4.1 considers the classical problem of optimizing a well-defined objective function. Section 6.4.2 focuses instead on optimizing performance in actual operations (the ultimate acid test).

A goal of this section is to bring out the different types of “optimality” that arise in stochastic, dynamic problems. This issue is rarely addressed in the context of static, deterministic problems where the objective function is well defined.

### 6.4.1 Calculating costs

Most optimization problems in a stochastic, dynamic setting would like to find a function  $x^\pi$  that minimizes (or maximizes, if  $c$  is a contribution) the following function:

$$F(\pi|\mathcal{T}_\infty, \Omega) = E \left\{ \sum_{t \in \mathcal{T}_\infty} \alpha^t c_t(x^\pi(I_t), K_t) \right\} \quad (8)$$

where  $0 \leq \alpha < 1$  is a discount factor and the expectation operator  $E$  is defined with respect to the probability space  $(\Omega, \mathcal{E}, \mathcal{P})_\infty$  defined over an infinite horizon  $\mathcal{T}_\infty$ . We include a discount factor which, combined with some standard assumptions on the boundedness of the problem, ensures that (8) remains finite. Note that  $K_t$  and  $I_t$  are both  $\mathcal{F}_t$ -measurable random variables.

For most modelers, this objective function is a kind of unachievable “ultimate” since, for hard problems, it is virtually impossible to prove infinite horizon optimality, and it is equally impossible to even compare two functions in this setting to see which is better.

Practical, experimental research must, as a rule, focus on calculating:

$$F(\pi|\mathcal{T}^{sh}, \Omega^{sh}) = \sum_{\omega \in \Omega^{sh}} p(\omega) \left\{ \sum_{t \in \mathcal{T}^{sh}} \alpha^t c_t(x^\pi(I_t), K_t) \right\} \quad (9)$$

This problem must be solved given the following equations:

Conservation of flow:

$$\sum_{d \in \mathcal{D}} x_{a,d,t} = R_{at} \quad \forall a \in \mathcal{A} \quad (10)$$

Rate of process transformation

$$x_{a,d,t} \leq u_{a,d,t} \quad (11)$$

System dynamics:

$$\begin{aligned} R_{a,t,t'}(\omega) &= R_{a,t-1,t'}(\omega) + \sum_{e \in \hat{\mathcal{E}}_t} \delta_{a,t,t'}(a'_e, d_{et}^0, \omega) x_{a'_e, d_{et}^0, t}^0 \\ &\quad + \sum_{d \in \mathcal{D}} \sum_{a' \in \mathcal{A}} \delta_{a,t,t'}(a', d_{t,t'}, \omega) x_{a', d, t} \quad a \in \mathcal{A}, \quad t' \geq t \end{aligned} \quad (12)$$

Information updating:

$$K_{t+1} = U^K(K_t, \kappa_t) \quad (13)$$

If forecasts are being updated based on simulated information (a process called “bootstrapping” which is not generally used in conventional modeling practice), we would include:

$$\rho_{t+1} = U^f(\rho_t, \kappa_t) \quad (14)$$

If we are using a value function, which is being updated dynamically based on sample information, we would include:

$$\hat{V}'_t \leftarrow U^V(\hat{V}_t, \hat{v}_t(\omega)) \quad (15)$$

If we have a multi-agent control structure, the problem is quite similar. We basically replace the time index  $t$  with the subproblem index  $q$  as follows:

$$F(\pi | \mathcal{T}^{sh}, \Omega^{sh}, \mathcal{Q}) = \sum_{\omega \in \Omega^{sh}} p(\omega) \sum_{q \in \mathcal{Q}} \{c_q(x_q^\pi(I_q), K_q)\} \quad (16)$$

#### 6.4.2 Measuring actual performance

The desire to optimize a well-defined, computable function in a controlled setting sometimes disguises the real problem, which is to change decisions as they are made in a real problem. Our efforts to optimize some function  $F_n$  must recognize that in most real applications, no-one actually cares what value we obtain for  $F_n$ . This function, ultimately, must be a surrogate for a more realistic problem.

Let  $x_t^a, t \in \mathcal{T}^{sh}$  represent a set of *actual* decisions made over a particular time period  $\mathcal{T}^{sh}$ . During this period of time, a single set of events  $\omega^a \in \Omega^a$  occurred. Given these events and decisions, we suggest that there are three major classes of objective functions that may be computed:

$$F(x^a | \mathcal{T}^{sh}, \omega^a) = \text{Surrogate cost function used by the optimization system. Here, we are using actual decisions, but evaluating them using our cost function.}$$

$$\begin{aligned}
G(x^a | \mathcal{T}^{sh}, \omega^a) &= \text{Management goal satisfaction - this function measures the degree to} \\
&\quad \text{which actual decisions produced results that match specific manage-} \\
&\quad \text{ment goals. These goals are specific, quantifiable measurements that} \\
&\quad \text{are used to evaluate system performance.} \\
H(x^\pi, x^a | \mathcal{T}^{sh}, \omega^a) &= \text{The user happiness function: this function provides a measure of how} \\
&\quad \text{well actual decisions } x^a \text{ match optimized decisions } x^\pi.
\end{aligned}$$

If we have developed a good optimization algorithm, we would expect to do well in terms of  $F(x^a)$ . However, if we do poorly in terms of the user happiness function  $H(x^\pi, x^a)$ , then we may not do well in terms of  $F(x^a)$ . Management, of course, wants to focus on their own measurements (which we call  $G$ ), and tends to be frustrated when modelers focus on  $F$  or users focus on  $H$ . The reality is that we have to work on all three. Our model only sees the function  $F$ , so this is the means by which we influence the behavior of the computer. We need to design algorithms that optimize  $F$  (if our algorithms do not optimize  $F$ , then changing  $F$  will have an unpredictable effect on system performance). We then need to work to adjust  $F$ , as well as train the users, to achieve the highest possible  $H$ , since a low value for  $H$ , by definition, means that the users are ignoring the model.

We can, of course, do well in terms of both  $F(x^a)$  and  $H(x^\pi, x^a)$  but poorly in terms of  $G(x^a)$ . If this happens, we have a case of a disconnect between user behavior and management goals. If we have indeed achieved a high level of user happiness (which means building trust and confidence in our model), then we have an opportunity to change actual decisions by slowly modifying the model.

## 7 Data representation

Using our mathematical model, we now have the elements needed to specify an instance of a DRiP in data. Below we list three types of information that comprise a DRiP: model related parameters and sets; raw data; and functions. We propose that this list, which we claim is comprehensive, can serve as a basis for a richer data specification that would allow modelers to share instances of relatively complex problems.

### Model specification:

$$\begin{aligned}
\mathcal{Q} &= \text{The set of subdivisions of the attribute space } \mathcal{A}. \\
\mathcal{A}_q &= \text{The set of elements of } \mathcal{A} \text{ in each subproblem.} \\
\mathcal{C}_q &= \text{The set of control classes for subproblem } q.
\end{aligned}$$

$$\begin{aligned}\mathcal{T}_q &= \text{The set of time elements for subproblem } q. \\ \mathcal{T}^{ih}, \mathcal{T}_t^{ph}, \mathcal{T}_t^{sh} &= \text{Implementation, planning and simulation horizons.}\end{aligned}$$

**Problem data:**

$$\begin{aligned}\mathcal{C}^I &= \text{The set of information classes} \\ \mathcal{C}^R &= \text{The set of resource classes} \\ \mathcal{C}^D &= \text{The set of decision classes} \\ \mathcal{R}_0^{(c)} &= \text{The set of resources in class } c \text{ at time period } 0, c \in \mathcal{C}^R. \\ \mathcal{L} &= \text{The set of resource layers (sets of resources that may be coupled together)} \\ a_r &= \text{Attribute vector of the } r^{\text{th}} \text{ resource in the first time period, } r \in \mathcal{R}_0^{(c)}, a_r \in \mathcal{A}^{(c)}. \\ \mathcal{D}^{(c,a)} &= \text{Set of possible decisions in class } c \in \mathcal{C}^D \text{ for resources in state } a \in \mathcal{A}. \\ K_0 &= \text{Initial data set of resources and parameters.} \\ \Omega^{ph} &= \text{Forecasted exogenous information (required if forecasting is involved).} \\ \rho_{0,t} &= \text{Initial set of forecast parameters for points in time } t \in \mathcal{T}^{ph} \text{ (generally required if forecasting is involved).} \\ u_{\mathcal{A},\mathcal{T}}(d, \omega) &= \text{Upper bounds on the rate of process transformation. We present this here as data, but it can also be a function (if applicable).} \\ x_{\mathcal{T}^{ph}}^p(\Omega^p) &= \text{A plan (or set of plans (if } |\Omega^p| > 1), \text{ which have been communicated to components that are exogenous to our system. Although rarely included in models, a “plan” can represent an important piece of data because of the cost of changing a plan (if applicable).}\end{aligned}$$

**Functions:**

$$\begin{aligned}M(a, d_t, K_t) &= \text{The modify function that determines how a resource is changed as a result of endogenous and exogenous factors. Implicit to the modify function are its component cost and transfer time functions.} \\ x_q^\pi(\mathcal{I}_q) &= \text{A function giving a decision (or set of decisions) } x_{dq}, d \in \mathcal{D}_q \text{ to be made at time } t \text{ given an } \textit{information set} \mathcal{I}_q \text{ for subproblem } q. \\ G^{(n)} &= \text{The attribute aggregation function for aggregation level } n \text{ (if applicable).} \\ \hat{V}_{t,t'}(R_{t,t'}) &= \text{The (approximate) value of having } R_{t,t'} \text{ resources at time } t' \text{ in a problem being solved at time } t \text{ (of course, this function may be zero) (if applicable).}\end{aligned}$$

- $U^K(K_t, \kappa_t)$  = The updating function for the database (required for any stochastic model).
- $U^J(\rho_t, \kappa_t)$  = The updating function for forecast parameters, given new information  $\kappa_t$  (if applicable - forecasts are generally fixed for a given model)
- $U^V(\hat{V}_t, \hat{v}_t)$  = The updating function for value function approximations (if applicable)

The modify function captures *all* of the physics of our problem. The decision function captures the intelligence. For many problems, the only dynamic data (within the model) is the set of resources, which is captured by the initial set  $\mathcal{R}_0$  and the updates in  $\Omega^{ph}$ .

## 8 Summary

This paper offers a new way of thinking about dynamic resource management problems, a problem class which we have named dynamic resource transformation problems. We have tried to offer as much generality as possible, subject to the constraint that the properties of the problem be sufficiently well defined that the problem can be easily represented as a well defined set of data and software. An important contribution of our representation is that it exposes many of the prominent dimensions of dynamic problems that are often lost or ignored in specific problems.

A software library built around this representation has been developed in Java. Elements of the architecture of this system are described in Shapiro & Powell (1998a).

We have intentionally left out any algorithmic details, since these are viewed as being problem specific. This decision, of course, raises the common question in operations research, Why model a problem that cannot be solved? A central claim of our representation is that any problem can be decomposed into suitably small subproblems and then solved. In Shapiro & Powell (1998b) we propose an algorithmic metastrategy that describes how DRiP's can be decomposed and solved using a class of approximation strategies drawn from dynamic programming.

## Acknowledgement

This research was supported in part by grant AFOSR-F49620-93-1-0098 from the Air Force Office of Scientific Research. We would also like to acknowledge the helpful suggestions of Teodor Crainic, as well as the students and staff of CASTLE Laboratory who struggled repeatedly through earlier

drafts of this document. I would also like to thank Erhan Cinlar for several valuable conversations that helped improve the notational style as well as the handling of the probabilistic aspects of information theory. In addition, the paper benefited from the careful comments of a conscientious referee.

## References

- Adelman, D., Nemhauser, G., Padron, M., Stubbs, R. & Pandit, R. (1996), Allocating fibers in cable manufacturing, Tech. Rep. LEC - 96 - 11, Georgia Institute of Technology.
- Andreatta, G. & Romanin-Jacur, G. (1987), 'The flow management problem', *Transportation Science* **21**(4), 249–253.
- Aronson, J. & Chen, B. (1989), A computational study of empirical decision horizons in infinite horizon multiperiod network flow models, Working paper 89-275, Dept. of Mgmt. Sci. and Information Tech., University of Georgia.
- Bertsekas, D. & Tsitsiklis, J. (1996), *Neuro-Dynamic Programming*, Athena Scientific, Belmont, MA.
- Bes, C. & Sethi, S. (1988), 'Concepts of forecast and decision horizons: Applications to dynamic stochastic optimization problems', *Mathematics of Operations Research* **13**(2), 295–310.
- Bhaskaran, S. & Sethi, S. (1987), 'Decision and forecast horizons in a stochastic environment : A survey', *Optimal Control Applications and Methods* **8**(1), 61–67.
- Birge, J. (1985), 'Decomposition and partitioning techniques for multistage stochastic linear programs', *Operations Research* **33**(5), 989–1007.
- Bodin, L. & Golden, B. (1981), 'Classification in vehicle routing and scheduling', *Networks* **11**, 97–108.
- Chen, Z.-L. & Powell, W. (to appear), 'A convergent cutting-plane and partial-sampling algorithm for multistage stochastic linear programs with recourse', *Journal of Optimization Theory and Applications*.
- Cheung, R. & Powell, W. (1996), 'An algorithm for multistage dynamic networks with random arc capacities, with an application to dynamic fleet management', *Operations Research* **44**(6), 951–963.
- Crainic, T. & Roy, J. (1991), 'Design of regular intercity driver routes for the LTL motor carrier industry', *Transportation Science*.
- Crainic, T., Gendreau, M. & Dejax, P. (1993), 'Dynamic stochastic models for the allocation of empty containers', *Operations Research* **41**, 102–126.
- Eiselt, H., Laporte, G. & Thisse, J.-F. (1993), 'Competitive location models: A framework and bibliography', *Transportation Science* **27**, 44–54.
- Fourer, R., Gay, D. & Kernighan, B. (1990), 'A mathematical programming language', *Management Science* **36**, 519–554.
- Geoffrion, A. (1989a), 'The formal aspects of structured modeling', *Operations Research* **37**(1), 30–51.
- Geoffrion, A. (1989b), 'An introduction to structured modeling', *Management Science* **33**, 547–588.

- Geoffrion, A. (1992a), ‘The SML language for structured modeling: Levels 1 and 2’, *Operations Research* **40**(1), 38–57.
- Geoffrion, A. (1992b), ‘The SML language for structured modeling: Levels 3 and 4’, *Operations Research* **40**(1), 58–75.
- Godfrey, G. A. & Powell, W. B. (1999), An approximation algorithm for solving multistage stochastic dynamic resource allocation problems, Technical Report SOR-99-XX, Department of Civil Engineering and Operations Research, Princeton University, Princeton, NJ.
- Gross, D. & Harris, C. (1985), *Fundamentals of Queueing Theory*, John Wiley and Sons, New York.
- Haghani, A. (1989), ‘Formulation and solution of a combined train routing and makeup, and empty car distribution model’, *Transportation Research* **23B**(6), 433–452.
- Hane, C., Barnhart, C., Johnson, E., Marsten, R., Nemhauser, G. & Sigismondi, G. (1995), ‘The fleet assignment problem: Solving a large-scale integer program’, *Math. Prog.* **70**, 211–232.
- Herren, H. (1973), ‘The distribution of empty wagons by means of computer: An analytical model for the Swiss Federal Railways (SSB)’, *Rail International* **4**(1), 1005–1010.
- Higle, J. & Sen, S. (1991), ‘Stochastic decomposition: An algorithm for two stage linear programs with recourse’, *Mathematics of Operations Research* **16**(3), 650–669.
- Infanger, G. & Morton, D. (1996), ‘Cut sharing for multistage stochastic linear programs with interstate dependency’, *Mathematical Programming* **75**, 241–256.
- Jordan, W. & Turnquist, M. (1983), ‘A stochastic dynamic network model for railroad car distribution’, *Transportation Science* **17**, 123–145.
- Mendiratta, V. & Turnquist, M. (1982), ‘A model for the management of empty freight cars’, *Trans. Res. Rec.* **838**, 50–55.
- Morton, D. (1996), ‘An enhanced decomposition algorithm for multistage stochastic hydroelectric scheduling’, *Annals of Operations Research* **64**, 211–235.
- Odoni, A. (1986), The flow management problem in air traffic control, in L. B. A.R. Odoni & G. Szego, eds, ‘Flow Control of Congested Networks’, Springer-Verlag, New York, pp. 269–288.
- Pereira, M. & Pinto, L. (1991), ‘Multistage stochastic optimization applied to energy planning’, *Mathematical Programming* **52**, 359–375.
- Pinedo, M. (1995), *Scheduling: Theory, Algorithm, and System*, Prentice Hall, Englewood Cliffs, NJ.
- Powell, W. (1986), ‘A local improvement heuristic for the design of less-than-truckload motor carrier networks’, *Transportation Science* **20**(4), 246–257.
- Powell, W. (1991), ‘Optimization models and algorithms: An emerging technology for the motor carrier industry’, *IEEE Transactions on Vehicular Technology* **40**, 68–80.
- Powell, W. (1996), ‘A stochastic formulation of the dynamic assignment problem, with an application to truckload motor carriers’, *Transportation Science* **30**(3), 195–219.
- Powell, W. & Carvalho, T. A. (1998), ‘Dynamic control of logistics queueing network for large-scale fleet management’, *Transportation Science* **32**(2), 90–109.
- Powell, W. B. & Shapiro, J. A. (1996), A dynamic programming approximation for ultra large-scale dynamic resource allocation problems, Technical Report SOR-96-06, Department of Civil Engineering and Operations Research, Princeton University.

- Puterman, M. L. (1994), *Markov Decision Processes*, John Wiley and Sons, Inc., New York.
- Schrijver, P. (1993), *Supporting Fleet Management by Mobile Communications*, P.R. Schrijver, The Netherlands.
- Sethi, S. & Bhaskaran, S. (1985), 'Conditions for the existence of decision horizons for discounted problems in a stochastic environment', *Operations Research Letters* **4**(2), 61–64.
- Shapiro, J. A. & Powell, W. B. (1998*a*), An object-oriented framework for dynamic resource transformation problems, Technical Report SOR-98-07, Department of Civil Engineering and Operations Research, Princeton University.
- Shapiro, J. A. & Powell, W. B. (1998*b*), A solution meta-strategy for dynamic resource transformation problems, Technical Report SOR-98-XX, Department of Civil Engineering and Operations Research, Princeton University.
- Sutton & Barto (1998), *Reinforcement Learning*.
- Van Slyke, R. & Wets, R. (1969), 'L-shaped linear programs with applications to optimal control and stochastic programming', *SIAM Journal of Applied Mathematics* **17**, 638–663.
- White, D. (1996), 'Decision roll and horizon roll processes in infinite horizon discounted Markov decision processes', *Management Science* **42**(1), 37–50.