SANDIA REPORT
SAND Number Pending
Unlimited Release (Pending R&A)
Printed September 2017

The CPAT 2.0.2 Domain Model - How CPAT 2.0.2 "Thinks" From an Analyst Perspective

Lucas A. Waddell, Frank M. Muldoon, Darryl J. Melander, Peter B. Backlund, Stephen M. Henry, Matthew J. Hoffman, April M. Zwerneman, Craig R. Lawton, Roy E. Rice

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology and Engineering Solutions of Sandia, LLC.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone:	(865) 576-8401
Facsimile:	(865) 576-5728
E-Mail:	reports@osti.gov
Online ordering:	http://www.osti.gov/scitech

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312

Telephone:	(800) 553-6847
Facsimile:	(703) 605-6900
E-Mail:	orders@ntis.gov
Online order:	http://www.ntis.gov/search

SAND # Pending
Printed September 2017
Unlimited Release (Pending R&A)

The CPAT 2.0.2 Domain Model - How CPAT 2.0.2 "Thinks" From an Analyst Perspective

[1]Lucas A. Waddell, [1]Frank M. Muldoon, [2]Darryl J. Melander, [3]Peter B. Backlund, [1]Stephen M. Henry, [4]Matthew J. Hoffman, [5]April M. Zwerneman, [5]Craig R. Lawton, [6]Roy E. Rice

[1]Systems Readiness and Sustainment Technologies, Org 8833
[2]Software Systems R&D, Org 9365
[3]Systems Technology, Org 5837
[4]Operations Research and Computational Analysis, Org 8831
[5]Mathematical Analysis and Decision Science, Org 8834

Sandia National Laboratories
P.O. Box 5800, Albuquerque, New Mexico 87185-1131

[6]Teledyne Brown Engineering
300 Sparkman Drive, Huntsville, AL 35805-1912

Abstract

To help effectively plan the management and modernization of their large and diverse fleets of vehicles, the Program Executive Office Ground Combat Systems (PEO GCS) and the Program Executive Office Combat Support and Combat Service Support (PEO CS&CSS) commissioned the development of a large-scale portfolio planning optimization tool. This software, the Capability Portfolio Analysis Tool (CPAT), creates a detailed schedule that optimally prioritizes the modernization or replacement of vehicles within the fleet - respecting numerous business rules associated with fleet structure, budgets, industrial base, research and testing, etc., while maximizing overall fleet performance through time. This report contains a description of the organizational fleet structure and a thorough explanation of the business rules that the CPAT formulation follows involving performance, scheduling, production, and budgets. This report, which is an update to the original CPAT domain model published in 2015 (SAND2015-4009), covers important new CPAT features.

This page intentionally left blank

TABLE OF CONTENTS

FIGURES

NOMENCLATURE

Abbreviation	Definition
CPAT	Capability Portfolio Analysis Tool
TWV	Tactical Wheeled Vehicles
FRP	full rate production
LRIP	low rate initial production
O&S	operations and sustainment
RDT&E	research, development, testing and evaluation
COA	course of action
MSR	minimum sustaining rate

This page intentionally left blank

PREFACE

This report serves as an update to the original Capability Portfolio Analysis Tool (CPAT) domain model, which was published in 2015[1]. All relevant content from the original publication remains in the interest of completeness. In addition, the following new CPAT features are covered:

- *Components* – A "component" is a division of the fleet into separate operational units with similar structure. See Section 2: Fleet Structure.
- *Mission and System Classes* – Mission and system "classes" are used to aggregate results to a level that provides meaningful, comprehensible information. Classes are especially useful in situations where a fleet has a large number of mission roles and/or system variants. See Section 2: Fleet Structure.
- *Batch Sizes* – Batch size refers to the smallest increment of a system that may be purchased. See Section 4.1: The Production Process.
- *Maximum Time New Systems are in Storage* – The analyst may specify an upper limit on the amount of time a new system may spend in storage before being fielded. See Section 4.2.2: Fielding New Systems from Storage and Business Rule 11.1.12: Fielding New Systems from Storage.
- *Minimum Cumulative Delivery* – The analyst may specify a minimum cumulative production for each product family. See Section 4.4.5: Product Family Minimum Cumulative Delivery and Business Rule 11.6.6: Family Minimum Cumulative Capacity.
- *Product Family Obviation* – The analyst may specify that a product family obviates one or more other product families. I.e., if and when a system from that product family is produced, systems from the obviated product families cannot be produced any longer. See Section 4.4.10: Product Family Obviation and Business Rule 11.6.13: Product Family Obviation.
- *Product Family Ratios* – The analyst may specify product family ratios which require that systems produced from a family must be distributed among components according to that ratio. See Section 4.4.11: Product Family Ratios and Business Rule 11.6.14: Product Family Ratios.
- *System Coasting* – The analyst may specify some systems as "coastable." A coastable system may be delivered into future time periods at the same rate that it was delivered in the final conventional time period. See Section 4.6: System Coasting and Business Rule 11.8: Coasting Systems.
- *RDT&E Active Costs* – Similar to procurement active costs, the analyst may specify an amount of RDT&E costs to be incurred in each period that a product family is active. This is in addition to any upfront RDT&E costs a product family may have. See Section

[1] Melander, Darryl J., Stephen M. Henry, Matthew J. Hoffman, Gio K. Kao, Craig R. Lawton, Frank M. Muldoon, Roy E. Rice, Liliana Shelton, "The CPAT Domain Model – How CPAT "Thinks" From an Analyst Perspective," Sandia National Laboratories, Albuquerque, NM SAND2015-4009.

5.3: Research Expenditures (RDT&E) and Business Rule 11.6.3: Family Per-Period Costs.

- *Component Earmarks* – Money may be allocated to purchases and upgrades of systems in a specific component (beyond the general combined expenditures budget). See Section 5.4.1: Component Earmarks and Business Rule 11.5.3: Component Earmarks.
- *Product Family Earmarks* – Money may be allocated to purchases and upgrades of systems in a specific product family (beyond the general combined expenditures budget). See Section 5.4.2: Product Family Earmarks and Business Rule 11.5.4: 52Product Family Earmarks.
- *Post-Processing Allocation of Expenditures* – For several CPAT results, it is necessary to allocate expenditures to various mission roles post-optimization. A new section has been added that covers the details of how the post-processing allocation is performed. See Section 5.5: Post-Processing Allocation of Expenditures to Mission Roles.
- *Economic Useful Life* – All systems in the fleet are given an economic useful life, which is the maximum age the system is allowed to be before it must be retired from a mission. See Section 6.8: Economic Useful Life and Business Rule 11.4.11: Economic Useful Life.
- *Arbitrary Phase Ordering* – The analyst may specify the optimization phase ordering to whatever suits their current analysis needs (e.g., schedule, age, yearly budget, horizon budget, performance, and cost). See Section 8: Constraint Violations and Intra-Tier Optimization P.
- *Post-Processing Calculation of Average Age* – Several CPAT results require the average age of systems of a given type to be calculated. A new section has been added that covers the details of how the post-processing age calculation is performed. See Section 10: Post-Processing Calculation of Average Age.
- *Optional Storage Upgrades* – The analyst has the option to specify for each upgrade whether that upgrade can be done in storage. See Business Rule 11.1.8: Optional Storage Upgrades.
- *Hard Limits on Tier Phases* – The analyst has the ability to enforce hard constraints on any of the phase objective metrics. See Section 8: Constraint Violations and Intra-Tier Optimization P and Business Rule 11.2.3: Hard Limits on Tier Phases.
- *Disallow Instantaneous Cross-Mission Transfers* – Systems retired in any time period cannot be immediately re-fielded in the same time period. See Business Rule 11.1.13: Disallow Instantaneous Cross-Mission Transfers.
- *Allow Overlap in System Obviations* – The analyst may specify a number of time periods in which both systems (the obviated and the obviating) can be delivered. See Section 4.5: System Obviation and Business Rule 11.4.7: System Obviation.

1. INTRODUCTION

Program executives face the perpetual fleet management challenge of devising investment strategies to assure optimal fleet modernization and to mitigate system obsolescence. These investment plans must be comprehensive, ensuring a balance between capability, schedule, and cost. This is particularly true for the United States Army where capability requirements must be met without violating increasingly strict expenditure limits, which are set in various categories including procurement, recapitalization, operations & support (O&S), and research, development, testing & evaluation (RDT&E).

The Capability Portfolio Analysis Tool (CPAT) supports fleet modernization planning by identifying a schedule of upgrades that yields the greatest overall fleet performance while adhering to budgets, production constraints, and other business rules. Leveraging a mathematical model that incorporates these concepts, CPAT helps decision-makers create and evaluate real-world fleet-wide modernization plans.

This document introduces the concepts, assumptions, and constraints built into CPAT and its underlying mathematical model. For more details and the actual mathematical constructs (variables, objectives, constraints, etc.) refer to (Waddell et al., 2017)[2]. For historical context regarding the CPAT project and its application and impact, see (Davis et al., 2016)[3].

[2] Waddell, Lucas A., Frank M. Muldoon, Stephen M. Henry, Matthew J. Hoffman, April M. Zwerneman, Peter B. Backlund, Darryl J. Melander, Craig R. Lawton, Roy E. Rice. “The Capability Portfolio Analysis Tool (CPAT): A Mixed Integer Linear Programming Formulation for Fleet Modernization Analysis (Version 2.0.2),” Sandia National Laboratories, Albuquerque, NM, SAND# Pending.

[3] Davis, Scott J., et al. “Maximizing the U.S. Army’s Future Contribution to Global Security Using the Capability Portfolio Analysis Tool (CPAT)”. *Interfaces*, 46.1 (2016): 91-108.

2. FLEET STRUCTURE

The systems in a fleet are organized into a hierarchical structure consisting of components, sets, groups, missions, and systems:

- A *component* is a division of the fleet into separate operational units with similar structure.
- A *set* is a collection of similarly configured groups.
- A *group* is a collection of resources working as a unit, such as a military brigade.
- Each group fulfills one or more *mission roles*. Each group in a set fulfills the same set of mission roles.
- Each mission requires a certain number of *systems* assigned to it in order to carry out that mission. The systems in a given mission role within a group may change over time, but in any given time period, each system occupies one and only one mission role within a group in a component.

This configuration is shown in Figure 1 below.

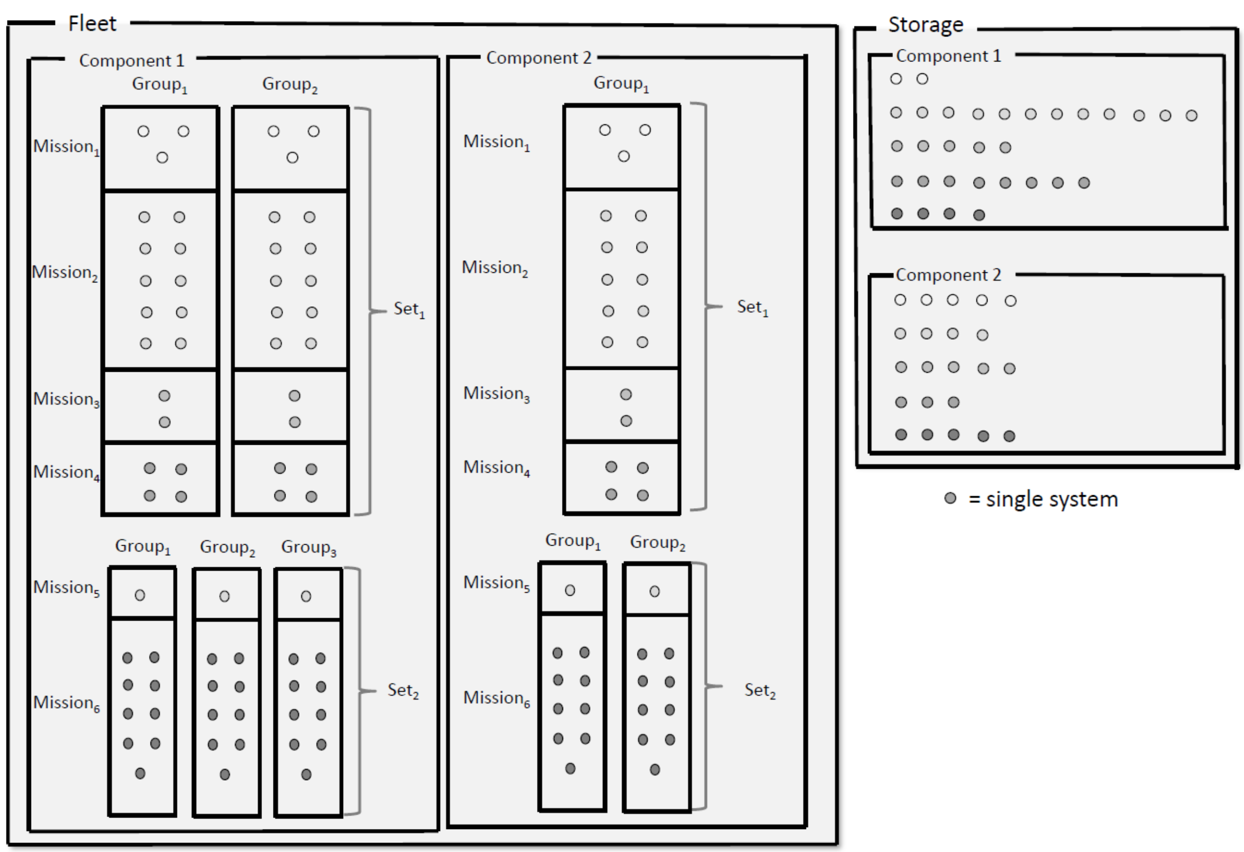

Figure 1. Fleet Structure

A fleet is described by defining each element of the fleet structure, including the components, the sets, the number of groups in each set, the missions each group must fulfill, and the number of systems required by each mission. The fleet description also includes the available system types, their properties, and the types of systems occupying each mission in each component at the beginning of the study horizon.

Often, the fleet is so large that it includes hundreds of unique system variants and mission roles. In these cases, certain optimization results provided by CPAT are so encyclopedic that they become incomprehensible to an analyst or decision maker. To alleviate this complexity, we employ the concepts of *mission classes* and *system classes* as methods of aggregating results (such as expenditures or performance improvement) to a level that provides meaningful yet comprehensible information.

Classes create a partitioning of the missions and systems across all components – each mission is assigned to exactly one mission class and each system is assigned to exactly one system class. These class partitions can be defined in any manner that is useful to the analyst and do not have to follow the fleet *component*, *set*, and *group* organization (i.e., two missions in different sets or components can still be in the same mission class). None of the business rules discussed later in this document act on the mission or system classes. In essence, these classes provide a "post-processing" organizational concept to aid understanding of enormous optimization plans.

A classic example of a mission class in the Tactical Wheeled Vehicles (TWV) fleet involves the "Light Tactical Vehicle" (LTV) missions. These missions 1) are quite diverse (e.g., Ambulance, Utility, Close Combat Weapons Carrier), 2) include a broad array of different systems (HMMWV variants), and 3) are divided into various sets (e.g., Armored Brigade Combat Team, Infantry Brigade Combat Team), but as a general organizational concept they all include vehicles in the fleet that are lighter in weight. A decision maker might naturally be interested in the proportion of budget spent overall on the LTV mission class vs, say, the Heavy Tactical Vehicle (HTV) mission class.

A standard example of a system class in the TWV fleet is the "Joint Light Tactical Vehicle" (JLTV) class. This collection of systems has several variants that are all produced by a common vendor, and thus many optimization results are intuitive when aggregated to this level.

3. FLEET PERFORMANCE AND TRANSFORMATION

3.1. Fleet Performance

CPAT supports fleet modernization planning by identifying a fleet transformation schedule that will yield the greatest overall fleet performance throughout the study horizon. The fleet transformation schedule – a list of suggested system upgrades and replacements and when they should take place – is based on information about the initial fleet configuration, the effectiveness of each system type when serving a particular mission, and budgetary and production constraints.

Fleet performance is calculated by examining the number and types of systems in the fleet and how those systems are allocated to missions and components. Each system type has a performance rating (composed of multiple performance categories) for each mission in which the system may serve. This performance rating is composed of one or more individual performance categories. Furthermore, each mission is assigned a relative importance; each system's performance rating is scaled by the mission's relative importance. The fleet performance for a particular time period is the sum of the scaled individual performance ratings for each system in service for a particular mission for all components at that time. The total fleet performance for the entire study horizon is the sum of the per-period fleet performance values. CPAT seeks to select a fleet transformation schedule which maximizes total fleet performance.

3.2. System Transitions

Performance is improved by changing which systems are in the fleet. Every change to the fleet is carried out via a transition. A transition can be either an *upgrade* (a modification made to a system already in the fleet) or an *exchange* (the removal of one system from the fleet and the introduction of an alternative system). To distinguish from other upgrades (discussed later) we will call upgrades within the fleet "in-mission upgrades". In all cases, a system transition is an event which changes the systems in service and causes corresponding changes to the fleet's cost and performance characteristics.

For each mission, transitions are limited to specific sets of valid before-and-after system type pairs. For example, analysts may declare that, in the context of a particular mission, system type A can be replaced with system type B, but disallow transitions from B to A or from A to C.

Transitions for missions in all components may not occur in such a way that multiple system types are serving the same mission in the same group; the systems serving each mission in a group must all be the same type of system. However, this does not preclude a mission from having differing system types across separate groups within the same component or in different components.

4. SYSTEM PRODUCTION

System production is the process of obtaining or creating systems that may be fielded to the fleet. Production supports fleet transformation by providing new systems that can be used to replace old systems. The production model incorporated into CPAT does not attempt to fully represent the details of real-world industrial base dynamics and constraints, but includes enough detail to address many common production and acquisition considerations.

4.1. The Production Process

When a system is upgraded or purchased for any mission in a component, the process takes time and money. The production process, which can be seen in Figure 2, consists of an administrative period, followed by a production period, followed by system delivery. The administrative period establishes when per-unit procurement charges are incurred; the production period represents time spent building the system; and delivery is when the system is made available for use. The production process can be tailored using the following parameters:

- Per-unit Procurement Cost – The cost incurred for each system purchased or upgraded
- Administrative Delay – The length of the administrative period, which is the number of time periods from when the per-unit cost is incurred to when production begins
- Production Delay – The length of the production period, which is the number of time periods a system is in production before it is delivered
- Long Lead Cost Ratio – The portion of the procurement cost which is incurred one time period before the administrative period begins.
- Batch size – the number of systems per "batch", the smallest increment of systems that can be purchased.

The analyst provides values for these parameters for each type of system that can be purchased, and for each type of upgrade that should be allowed, as defined by seed (pre-upgrade) and target (post-upgrade) system type pairs.

The example in Figure 2 demonstrates an administrative delay of 2 time periods and a production delay of 2 time periods. Note that the administrative period always immediately precedes the production period.

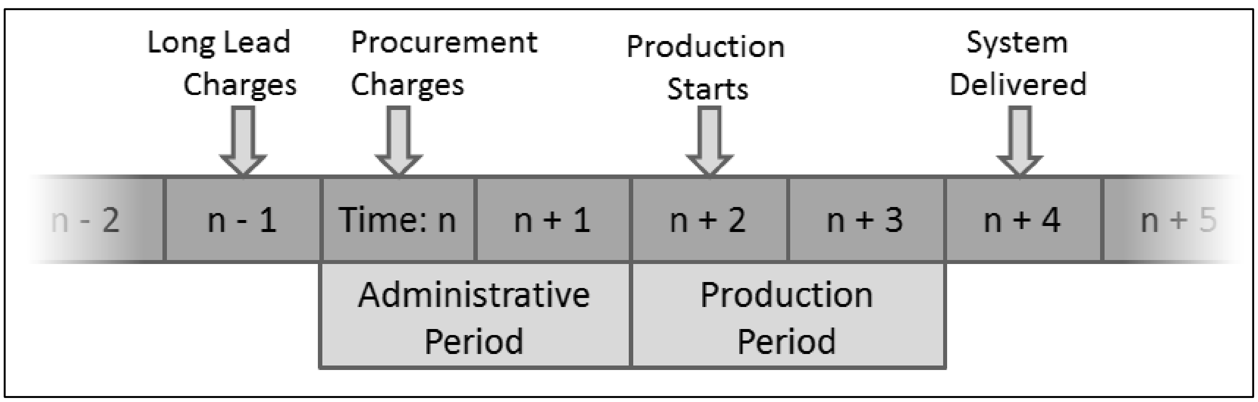

Figure 2. The Production Process

4.2. The Storage Yard

Systems removed from service in one mission in one component may be repurposed for use in another mission in the same component. This reapplication process is facilitated by sending systems to a conceptual storage yard, divided into components, where they are kept until fielded to another mission in the same component. Systems are available to be repurposed immediately; a system can be placed in a new mission in a component in the same time period it is removed from its former mission in the same component.

Furthermore, purchases can be made before the new systems are to be fielded to missions in a component. Pre-purchased systems are also held in a particular component in the storage yard until they are put into service.

Systems can be modified while in the storage yard. An *in-storage upgrade* represents modifications needed to prepare the system for its new mission. In-storage upgrades may only take place for those seed-target system type pairs that have been marked as valid by the analyst.

Recall that fielding a system from the storage yard is known as an *exchange*, so called because the systems being fielded from storage are traded with the old systems they are replacing. The number of systems sent to storage due to an exchange is always equal to the number of systems taken out of storage.

Note that system exchanges cause the old system to be sent to a component in the storage yard, but in-mission upgrades do not. Also be aware that a transition that may be executed as an in-mission upgrade may also be executed as an exchange, subject to system availability. If the transition is executed as an exchange, the old system is sent to a component in the storage yard even if an upgrade between the same two system types is also possible.

In-storage upgrades follow the same production process as other transitions, including an administrative period, a production period, and a per-system upgrade cost. The production period for an in-storage upgrade may begin as soon as the system is sent to the storage yard, but not before. The administrative period for an in-storage upgrade may begin before the system is sent to storage, but the system must arrive at the storage yard by the time the production period begins and remain in storage until production of the upgrade has been completed.

4.2.1. Storage Upgrade Consumption Priority

It is possible for multiple different seed system types to be upgraded to the same "target" system type in the storage yard. It may be desirable to create a priority ordering of these seed systems to determine which are used first to upgrade to the new target systems. The analyst can choose to enforce a *storage upgrade consumption priority* among these seed systems. A *storage upgrade consumption priority* determines the sequence in which seed systems in each component in storage are upgraded to any target system.

> Example: Assume that the optimization chooses to upgrade to System A in Component 1. System A is obtained via an in-storage upgrade from System B or System C. The analyst selects a consumption priority of System B before System C. CPAT must upgrade to System A in Component from System B while there are System B's available in storage in Component 1 before upgrades from System C are made.

4.2.2. *Fielding New Systems from Storage*

New systems delivered to storage via purchases or in storage upgrades may only remain in storage for a user defined number of time periods. This prevents systems from building up in storage where they do not age and then being fielded in a later time frame with an age of zero. A slack is allowed to be in storage indefinitely and not field. This slack is defined as one less than the maximum of a purchase batch size and the largest requirement for a mission that the system supports. This allows purchases for missions in which the purchased system batch size is not a divisor of the number of systems required for the mission. It also allows for situations where it may take several years of production to acquire enough systems to field an entire group's worth. Additionally, any LRIP systems or systems in initial storage may only remain in storage for the user-specified number of time periods, beginning in the first time period that non-LRIP production is completed.

Example: Assume that the optimization chooses to upgrade to System A in Mission M1 in Component 1 (and Mission M1 is the only mission that System A supports). M1 has four groups with 10 systems in each group. There are 17 units of System A initially in storage in Component 1, and System A is only acquired by purchases with a batch size of 8. The number of time periods newly purchased or in storage upgraded systems can stay in storage is 2. The optimization chooses to purchase 2 batches of System A in time period 3 and 1 batch of System A in time period 7. Three groups in Mission M1 in Component 1 must be modernized to System A by time period 5. The other group in Mission M1 in Component 1 must be modernized to System A in time period 7, time period 8, or time period 9.

4.3. Upgrades Trump Purchases

New systems can be obtained via purchases or upgrades. Some systems may be obtainable via both a new purchase and an upgrade from a seed system option. The analyst can choose to enforce that seed systems must be used, if available in that component, to upgrade to target systems before any purchases of those target systems are made.

Example: Assume that the optimization chooses to upgrade to System A in Component 1. System A is available as a purchase or an upgrade from System B. The analyst selects that System B must be used as an upgrade, if available, to any target system before purchases of that target system are made. CPAT must upgrade to System A from any System B's that are available in storage in Component 1 before any new purchases of System A are made.

4.4. Product Families

Some constraints and expenditures apply to collections of system types, such as all systems produced at a common facility. These are modeled as product families. A *product family* is simply a collection of system types. Each system type may be assigned to any number of families.

4.4.1. *Delivery Gaps*

For each product family, the analyst can choose whether to allow delivery gaps within the family. If gaps are not allowed, then delivery of the family's systems cannot start and then stop and then start again. This means that once any system in a family is delivered, at least one system in the family must be delivered each time period until the family becomes inactive. After the family first delivers a system, if there is ever a time period when no system in the family is delivered, then no system in the family can ever be delivered again.

Also be aware that fielding an existing system from storage does not require production; purchases and upgrades do. Therefore, systems in inactive families cannot be purchased, nor can they be the product of an in-mission or in-storage upgrade, but they can be the *input* to an upgrade, and may still be fielded from storage if the system is already being held there.

4.4.2. Product Family Minimum Sustaining Rates

It may not be feasible or desirable to deliver too few systems from a product family. For example, a product family may represent a production facility with a minimum sustaining rate (MSR). To prevent unreasonably small production batches, the analyst may specify a minimum delivery rate for each product family. For each time period in which at least one system from the family is delivered, there must be at least the minimum number of systems delivered. The family's final delivery year, however, is allowed to fall below the minimum delivery rate.

Example: A product family has a minimum delivery rate of 20 systems. In any given time period except for the family's final delivery year, CPAT may choose to deliver no systems from this family, or may choose to deliver 20 or more systems. CPAT may not choose to deliver between 1 and 19 systems except in the family's final delivery year.

4.4.3. Product Family Minimum Per-Period Delivery

It may not be feasible or desirable to deliver too few systems from a product family in a specific time period. For example, a product family may represent a production facility that has a minimum number of total systems that must be produced each time period to make the production economical or current plans might require specific quantities of systems from product families in certain time periods. To prevent unreasonably small production, the analyst may specify a minimum production for each product family for each time period. The product family must deliver the minimum amount of systems desired for each time period.

Example: A product family has a minimum delivery of 200 systems in time period 2. CPAT must deliver at least 200 systems from this family in time period 2.

4.4.4. Product Family Maximum Per-Period Delivery

It may not be feasible or desirable to deliver too many systems from a product family in a specific time period. For example, a product family may represent a production facility that has a maximum number of total systems that must be produced each time period. To prevent unreasonably high production, the analyst may specify a maximum production for each product family for each time period. The product family cannot deliver more than the maximum amount of systems desired for each time period.

Example: A product family has a maximum delivery of 200 systems in time period 2. CPAT must not deliver more than 200 systems from this family in time period 2.

4.4.5. Product Family Minimum Cumulative Delivery

It may not be feasible or desirable to produce too few cumulative systems from a product family. For example, a product family may represent a production facility that has a minimum number of total systems that must be produced to make the production economical. To prevent unreasonably small cumulative production, the analyst may specify a cumulative minimum production for each product family. If the product family produces any systems then the cumulative number produced must be at least the minimum number desired.

> Example: A product family has a cumulative minimum production of 200 systems. Over the planning horizon, CPAT may choose to produce no systems from this family, or may choose to produce 200 or more systems. CPAT may not choose to produce between 1 and 199 systems.

4.4.6. Product Family Maximum Cumulative Delivery

It may not be feasible or desirable to produce too many cumulative systems from a product family. For example, a product family may represent a production facility that has a maximum number of total systems that it is able to produce. To prevent too large of a cumulative production, the analyst may specify a cumulative maximum production for each product family. If the product family produces any systems then the cumulative number produced must be no greater than the maximum number desired.

> Example: A product family has a cumulative maximum production of 200 systems. Over the planning horizon, CPAT may choose to produce up to 200 systems, but no more, from the product family.

4.4.7. Production Smoothing

It may be desirable to prevent delivery rates from varying too dramatically from one year to the next. For each product family, the analyst may specify a maximum delivery variance. If a delivery variance is specified, then the most desirable median production level is determined by the optimization, and then yearly delivery from the family must fall within a band centered on that median level. The bandwidth is exactly determined by the specified delivery variance. The family's final delivery year, however, is allowed to fall beneath the delivery band as the production line winds down. There may also be a ramp-up period prior to full-rate production, during which delivery output is not required to respect production smoothing. Instead, the number of systems delivered must simply be non-decreasing in time during this ramp-up.

> Example: A product family has a delivery variance of 0.2. If the optimization determines that the most desirable median production level is 100 systems per year, then the number of systems delivered from the family in any time period must fall between 90 and 110 systems. During the final delivery year, the number of systems delivered is allowed to fall below 90 systems. If a ramp-up period of 2 years is defined, the first 2 years of production can fall below 90 systems, as long as the output in the second year is no less than the first.

4.4.8. Product Family Expenditures

Product families incur several types of expenditures: startup costs, per-period procurement costs, per-period RDT&E costs, and upfront RDT&E costs. The first three types are based on the product family's activity status, described below; upfront RDT&E costs are discussed in a later section.

Each product family is considered either active or inactive during each time period. A family is active whenever any of its systems (including LRIP) is in an administrative or production period. A cost per active period (procurement and/or RDT&E) is incurred for each time period the product family is active. A product family may also incur startup charges when any full-rate production (FRP) systems in the family first become active. Startup charges are entered as a list of amounts to be charged, and the number of time periods before or after an FRP system in the associated product family first becomes active that each charge should be incurred.

If a system belongs to more than one family, then producing a system causes all of its families to become active, and counts against all of its families' capacity limits and gap constraints.

4.4.9. LRIP

Low Rate Initial Production (LRIP) refers to systems that are produced before a product family enters its normal production state. It represents low-level production that takes place as production capacity is ramping up to normal levels. Some of these initial systems may be made available for fielding to missions in certain components, while others will serve other purposes and can never be used by missions in any component.

The analyst specifies the LRIP schedule for each system type in each product family. The LRIP schedule consists of the number and types of systems that will be produced in each of the five years before normal production begins, and the number of these systems that will be made available for use by missions in each of these years. The schedule is relative to the first system delivered due to normal production – the last LRIP system is delivered one year before the family's first non-LRIP system.

LRIP can be purchases or upgrades. LRIP incurs the same cost and takes the same amount of time as a normal purchase or upgrade. LRIP incurs the same per-period product family expenditures as normal production, and will cause per-period procurement and RDT&E expenditures to be incurred for *all* families having the LRIP system as a member. The differences between normal production and LRIP production are:

- The analyst specifies exactly how many systems of each type will be produced via purchase or upgrade
- LRIP is not subject to capacity schedules or minimum delivery rates
- Like normal production, LRIP production may "consume" the seed systems during an upgrade, in which case the required number of seed systems must be available in storage for LRIP to take place. However, for LRIP in particular the analyst may indicate that an LRIP upgrade should *not* consume seed systems, thus the seed systems do not need to be available in storage in order for the LRIP upgrades to take place. This is useful in cases where the correct upgrade costs and delays need to be incurred but the seed systems for LRIP need not (or in some rare cases *should* not) be explicitly represented in storage.

The optimization chooses which component LRIP systems are produced for and delivered to. Seed systems must come from the same component in which the LRIP systems are produced for. LRIP only occurs for those system types that are later delivered due to normal production. If no system of a particular type is ever produced, then that system's LRIP schedule will not occur.

4.4.10. Product Family Obviation

Sometimes production from product families will make production from other product families obsolete. Once any system type from a product family starts production, system types from the other product family are no longer able to be produced, possibly because production facilities have been reconfigured to support certain product families. CPAT supports this type of relationship between product families through *product family obviation*. Each product family may be obviated by any number of other product families. Once any system of these obviating product families is delivered, then any system in the obviated product family cannot ever be delivered again.

Example: Product family A is obviated by product family B and product family C. If any systems belonging to product family B are first delivered in 2015, and any systems belonging to product family C are first delivered in 2016, the latest that any system from product family A could be delivered is 2014, the time period just before any of the systems from the obviating product family are first delivered.

Product family obviation can also be used to force CPAT to choose to field only systems from at most one of two product families. This is done by making each product family obviate the other product family. More generally, multiple product family obviation constraints can be used collectively to force CPAT to choose to field only systems from at most one of several product families. This is done by making each product family obviate all of the other product families in the set from which only one should be fielded.

Example: The analyst wants CPAT to choose to field either systems from product family A or systems from product family B. The analyst indicates that product family A obviates product family B, and product family B obviates product family A. If CPAT chooses to produce systems from product family A, systems from product family B will never be produced. Similarly, if systems from product family B are ever produced, systems from product family A will never be produced. CPAT can choose to produce systems from one or the other, but not both.

4.4.11. Product Family Ratios

Sometimes it is desired for production of systems from a product family to be distributed among the components in the fleet. This prevents the possibility that one component could receive an excessive amount of modernized systems while other components receive none. The analyst can designate product family ratios which require that systems produced from a product family must be delivered to components by a specified ratio. Along with the ratios for the product family, the analyst specifies a variance and a delivery window. The variance specifies how far above or below the actual delivery of systems to each component can be from the specified ratio. The delivery window designates over how many time periods the product family delivery ratios are enforced.

Example: The analyst wants CPAT to field systems from Product Family A to Components 1 and 2 using a ratio of 8 to 2. The analyst specifies a delivery variance of 0.25 and a window of two time periods. For every 1 system delivered to Component 2 over two time periods from Product Family A, then between $\frac{8}{2} * (1 - 0.25) = 3$ and $\frac{8}{2} * (1 + 0.25) = 5$ systems must be delivered to Component 1 over the same two time periods.

If delivery to one component from a product family begins, then delivery to all components must also begin in the same time period. Once missions in a component have completely modernized to all systems from a product family, then deliveries to that component drop to zero and the ratio for that component is no longer enforced. All missions in other components still receiving deliveries from a product family must maintain the specified ratios until those missions in those components have completely modernized to systems from the product family.

4.5. System Obviation

Sometimes a new system type will make certain previously available system types obsolete. Once the new system type starts production, older system types are no longer able to be produced, possibly because production facilities have been reconfigured to support the new system type. CPAT supports this type of relationship between system types through *system obviation*. Each system type may be obviated by any number of other system types. Once any of these obviating systems is delivered, the obviated system cannot ever be delivered again.

Example: System A is obviated by System B and System C. If System B is first delivered in 2015, and System C is first delivered in 2016, the latest that System A could be delivered is 2014, the time period just before any of its obviating systems is first delivered.

In some cases, there may be a number of overlap time periods in which both systems (the obviated and the obviating) can be delivered. This is done by specifying a non-zero amount of overlap allowed for the system obviation pair. Overlap time periods start in the first time period in which the obviating system is delivered. The default overlap is 0.

Example: System A is obviated by System B, with 2 years of overlap. If System B is first delivered in 2015, the latest that System A could be delivered is 2016. This allows for 2 time periods of overlapping delivery with the obviating system.

System obviation can also be used to force CPAT to choose at most one type of system from a set of systems. This is done by marking each system in the set as obviating all other systems in the group.

Example: The analyst wants CPAT to choose between System A and System B. The analyst indicates that System A obviates System B, and System B obviates System A. If CPAT chooses to produce System A, System B will never be produced. Similarly, if System B is ever produced, System A will never be produced. CPAT can choose to produce one or the other, but not both.

4.6. System Coasting

Sometimes it is desired, for model simplification purposes, to reduce the fidelity of delivery decisions in the extended time horizon. As a way to do this, some systems are designated as *coastable systems*. That is, they are allowed to continue to deliver to all missions in the extended time horizon at the same rate as was delivered to those missions in the final conventional time period. Once a mission contains only coastable systems in the extended time horizon, then coasting to that mission stops. Delivery of coastable systems is allowed to ramp-down in the final delivery period to each mission in which a coastable system is fielding to. By convention, a mission will never be supported by both future systems and systems that are allowed to coast.

Example: System A is designated as a system that can coast into extended time periods. System A supports Missions 1 and 2. Mission 1 and Mission 2 both have 10 groups. In the final conventional time period, System A produces and fields 2 groups to Mission 1 and 0 groups to Mission 2. After the final conventional time period, Mission 1 has 5 groups of System A. In the extended time horizon, System A can continue to be produced and fielded at the same rate as it was in the last conventional time period at two groups per year until all 10 groups of Mission 1 are modernized to System A. Mission 2 is not allowed to modernize any more groups of System A after the final conventional time horizon because System A was not delivered to Mission 2 in the last conventional time period.

5. BUDGETS AND EXPENDITURES

CPAT tracks three types of expenditures: procurement, operations, and research. Each expense category has its own budget and its own set of charges. There is also a combined budget which limits the money available to an analyst-specified combination of the three budget categories. To briefly summarize the expense categories:

- Procurement expenditures represent the cost of system production, including upgrades, new purchases, and product family start-up and per-period procurement costs.
- Operating expenses represent the cost of using and maintaining systems present in the fleet.
- Research expenses represent the cost of any upfront prerequisite activities that make new system types eligible for procurement, as well as any research expenses incurred per-period by a product family.

For each of these expense categories, including the combined category, analysts can specify both a budget for each time period, and a horizon budget which limits the total expenditures in the category across all time periods in the study.

5.1. Procurement Expenditures

Procurement expenditures represent the cost of modernizing systems. Whenever one or more system is purchased or upgraded, an associated per-unit procurement cost is incurred. The cost of a purchase depends on the type of system purchased; the cost of an upgrade depends on the type of system before and after the upgrade.

Per-unit procurement costs are normally charged at the beginning of the procurement event's administrative period, *before* the new system is delivered.

> Example: Assume that a mission with 30 systems will be upgrading from System A to System B in 2015. If the A->B procurement cost is $100,000 per system, the A->B administrative delay is 0, and the A->B production delay is 2 years, then the mission will be charged $3,000,000 in 2013 (30 x $100k, 2 years before the delivery date).

The analyst may choose to incur a portion of the procurement costs a year early, one year before the administration period begins. These early charges, called long lead costs, can range from 0% to 100% of the procurement cost (Figure 2).

Some expenditures related to product families are also considered procurement expenditures. Product family startup costs, triggered when an FRP system in the family first becomes active, are incurred as procurement expenditures. Each period a family is active, the family's per-period procurement active cost is also incurred as a procurement expenditure.

For reporting purposes, product family expenditures are allocated to missions using a system density weighting method; the portion of each product family expenditure incurred by a given mission is proportional to the percentage of systems delivered from the product family that were fielded to that mission. Percentages are based on totals for the entire study horizon, not just the time period that the cost was incurred. This means that some charges may be incurred in time periods when the mission was not actively fielding systems.

Example: 30% of the systems delivered from product family PF1 were delivered to mission M. All of the systems delivered to mission M were produced in time period 5. Systems were produced from PF1 for other missions in time periods 10 and 11. Mission M incurs 30% of PF1's startup costs, and 30% of the per-period costs incurred in time periods 5, 10, and 11.

5.2. Operating Expenditures (O&S)

Operating expenditures represent the cost of operating and maintaining the systems in the fleet. Every time period, operating expenses are incurred for each system in service at that time. The operating cost per time period depends on the system type and the mission the system is serving. The analyst can specify a different operating cost for each mission that a particular system type may serve.

Example: Mission M1 has 30 systems. For the first 3 time periods, the mission uses system A. In time period 4 the mission switches to system B and continues to use system B for the 7 remaining periods of the 10 time period study horizon. If System A in M1 has an O&S cost of $10k/year, and System B in M1 has an O&S cost of $8k/year, the O&S cost for this mission is $300k/year for the first three years, and $240k/year for the next 7 years, for a total of $2,580,000.

5.3. Research Expenditures (RDT&E)

Some systems require upfront research before they become available. Similarly, some product families require research during time periods in which they are active (see Section 4.4.8). CPAT supports features which capture the costs of both these upfront and active period activities.

Each set of research activities is represented as a Research, Development, Testing and Evaluation (RDT&E) Effort. Each RDT&E Effort is associated with one or more system types, via assignment to a product family. If the fleet transformation schedule ever includes a system associated with the product family, then that family's RDT&E Effort upfront costs are incurred. If no system associated with the family is ever in the fleet, the effort's upfront costs are not incurred.

If the RDT&E Effort includes per-period costs, the family's RDT&E active cost is incurred in each period the family is active. If the family is never active, the effort's active costs are not incurred.

An RDT&E Effort's upfront costs and their timing are dependent on the family's delay – the difference between when a family's systems were first delivered, compared to when they potentially could have first been delivered. Put another way, the delay for a family is the smallest delay of any of its associated systems.

Each RDT&E Effort has one or more associated upfront cost profiles. Each profile is associated with a particular delay, and specifies the amounts and time periods when upfront research expenditures will occur (see Figure 3). The profile that matches the family's delay is the one that will be charged. By default, a family which has an RDT&E effort defined may not have a delay for which there is no cost profile (e.g., it cannot be delayed by 4 years if cost profiles are only defined for delays up to 3 years). An analyst may use this to restrict the maximum delay for a particular family. An RDT&E cost profile cannot be incurred if any of the costs occur in the extended time horizon.

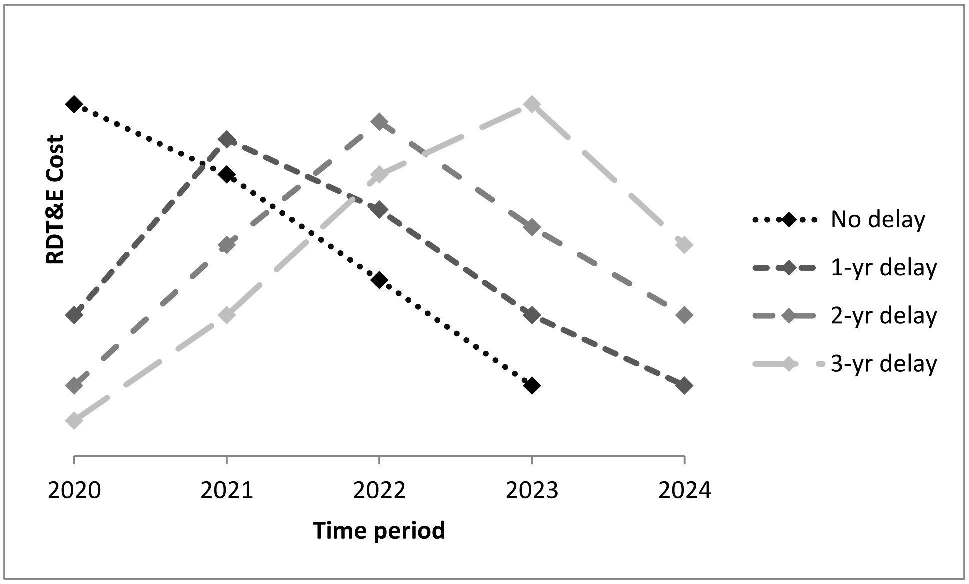

Figure 3. Notional Upfront RDT&E Cost Profile Example

The relationship between cost profiles and family delays can be relaxed by deselecting the "Enable RDT&E Expenditure Offsets" option. When this option is turned off, delays are not restricted to those for which there is a cost profile defined. If the RDT&E Effort is ever activated, the cost profile for zero delay is charged no matter what the family's actual delay happens to be. This is solely for backwards compatibility with models from earlier formulation versions.

Example: An RDT&E upfront cost profile has been defined for a product family, with System Z as the only system in that family (and therefore the only one requiring the RDT&E profile). The first period when System Z could be delivered is year 7. However, CPAT chooses to wait until year 9, two years later than it first becomes available. Because System Z was delayed by two periods, the delay for this family is 2. If the "Enable RDT&E Offsets" option is set, the RDT&E cost profile for a delay of 2 will be charged. If the option is not set, the RDT&E cost profile for a delay of zero will be charged. If system Z had never been produced, then no associated RDT&E costs would ever be charged.

RDT&E costs are distributed to missions based on system density weighting (see Section 5.5).

5.4. Combined Expenditures

Combined expenditures consist of an analyst-specified combination of any of the three individual expense categories. Setting a combined budget can be useful if the analyst wants to limit the total amount spent, but is not worried with how expenses are split among the categories. A combined budget can also be useful when used with individual category budgets. By setting the combined budget to less than the sum of the individual categories, the analyst can allow some flexibility in how expenditures are divided among expense categories, but still set specific limits for individual categories.

5.4.1. *Component Earmarks*

Additional money, above the per-period combined expenditures budget, may be allocated to purchases and upgrades of systems in a specific component in a specific conventional or extended time period. An earmark for a component does not require purchases and upgrades of at least that amount to be spent for that component, but it does not allow a portion of the earmark to be spent for the purchases and upgrades in another component.

Example: A $100 earmark for Component 1 is designated in time period 3. The Procurement Budget for the fleet in time period 3 is $500. At most $600 can be spent on upgrades and purchases of systems in Component 1 in time period 3. At most $500 can be spent on upgrades and purchases for systems in components other than Component 1 in time period 3. It would be allowable for $150 to be spent for purchases and upgrades to Component 1 and $450 to be spent for purchases and upgrades to other components in time period 3.

5.4.2. *Product Family Earmarks*

Additional money, above the per-period combined expenditures budget, may be allocated to purchases, upgrades, active costs, startup costs, LRIP, and RDT&E of systems in a specific product family in a specific conventional or extended time period. An earmark for a product family does not require purchases, upgrades, actives costs, startup costs, LRIP, and RDT&E of at least that amount to be spent for that product family, but it does not allow a portion of the earmark to be spent for purchases, upgrades, actives costs, startup costs, LRIP, and RDT&E in another part of the fleet. When a system is purchased or upgraded in a time period in which it could charge either a product family earmark or a component earmark, the product family earmark is charged.

Example: A $100 earmark for Product Family 1 is designated in time period 3. The Procurement Budget for the fleet in time period 3 is $500. At most $600 can be spent on purchases, upgrades, actives costs, startup costs, LRIP, and RDT&E of systems in Product Family 1 in time period 3. At most $500 can be spent on purchases, upgrades, actives costs, startup costs, and RDT&E for systems not in Product Family 1 in time period 3. It would be allowable for $150 to be spent for purchases, upgrades, actives costs, startup costs, LRIP, and RDT&E to Product Family 1 and $450 to be spent on purchases, upgrades, actives costs, startup costs, LRIP, and RDT&E in other parts of the fleet in time period 3.

5.5. Post-Processing Allocation of Expenditures to Mission Roles

For several CPAT results, it is necessary to perform a post-optimization allocation of all expenditures to the various mission roles within the fleet. Each mission role (combination of system, mission, and component) is assigned a share of the overall fleet expenditures in each time period based on a number of different rules.

1. Some costs are already incurred within the context of a mission role, so no adjustment to allocation is needed. This includes the costs associated with *in-mission upgrades, coasting system purchases, coasting system upgrades, future system transitions, and*

future system LRIP. CPAT knows the number of systems fulfilling each mission in each component in each time period, so *O&S costs* are also easily processed in this manner.

2. Other costs are incurred for a particular system within the context of a component, but not for any one given mission. This includes the costs associated with *purchases* and *in-storage upgrades*. Over the entire time horizon, a certain proportion of each system in each component fielded from storage will support each mission. A mission gets assigned a share of purchase or in-storage upgrade costs based on this proportion.
   In the special case that *no* systems of a certain type are ever fielded from storage for a component (e.g., only via in-mission upgrades), the proportion is based on *total* fielding. This case only occurs when the solution is not fully optimal and includes unnecessary purchases.

Example: Suppose that there are three missions, X, Y, and Z. In time period 2, $500 is spent on purchases of System A in Component 1, and $100 is spent on in-storage upgrades to System B in Component 1. Over the entire time horizon, the total number of Systems A and B fielded to Missions X, Y, and Z in Component 1 is given by the following table. The numbers in parentheses represent the number of systems fielded *from storage*.

	X	Y	Z
A	1000(800)	700(200)	300(0)
B	200 (0)	200 (0)	100 (0)

The $500 purchase cost for System A in Component 1 in time period 2 is allocated in the following manner (based on fielding *from storage*): $400 to Mission X, $100 to Mission Y, $0 to Mission Z. The $100 in-storage upgrade cost for System B in Component 1 in time period 2 is allocated in the following manner (based on total fielding, since there is never any fielding from storage): $40 to Mission X, $40 to Mission Y, $20 to Mission Z.

3. *LRIP production costs* are slightly different in that they are not incurred within the context of a component. These costs are allocated to missions in exactly the same way as purchase and in-storage upgrade costs, with the exception that the proportions are based on fielding across the *entire fleet*, not just within a certain component.

4. *Product family upfront RDT&E costs* are allocated to mission roles based on the proportion of total systems fielded from the product family (across the entire time horizon) that support that mission role.

Example: Over the entire time horizon, 1000 systems from Product Family PF1 are fielded. Of these, 200 are System A fielded to Mission X in Component 1. Therefore 20% of all PF1 upfront RDT&E costs will be assigned to this mission role.

5. In a given time period, *product family procurement active costs* and *product family RDT&E active costs* are allocated only to mission roles whose systems are active (in an administrative or production period) in that time period. The costs are assigned to system types based on the proportion of the purchase costs of all active systems that are incurred for that type. Then, for each system type, the costs are allocated to mission roles based on the proportion of total systems of that type fielded to that mission role.

Example: Product Family PF1 consists of systems A, B and C, and has a per-period procurement active cost of $160. System A supports missions M1 and M2, System B supports missions M1 and M2, and System C supports Mission M2. The following table gives the quantities of systems A, B, and C fielded throughout the entire time horizon.

	Component 1		**Component 2**	
	M1	**M2**	**M1**	**M2**
A	2000	1000	200	800
B	200	200	800	800
C	0	500	0	450

In time period 3, 200 System A's are in production, and 300 System B's are in an administrative period. Each System A costs $10 to purchase, while each system B costs $20 to purchase. The $160 procurement active cost for Product Family PF1 in time period 3 gets allocated to mission roles as follows.

	Component 1		**Component 2**	
	M1	**M2**	**M1**	**M2**
A	$20	$10	$2	$8
B	$12	$12	$48	$48
C	$0	$0	$0	$0

Mission roles involving System C are not assigned any portion of the procurement active cost in time period 3 because no system C's are active.

Note: Per-period RDT&E costs are allocated to mission roles in the same manner.

6. *Product family startup costs* are allocated only to mission roles whose systems belong to that product family. The costs are assigned to system types based on each system's proportion of the cumulative purchase and upgrade costs of all systems in the product family. Then, for each system type, the costs are allocated to mission roles based on the proportion of total systems of that type fielded to that mission role.

Example: Consider the horizon fielding schedule for the systems in PF1 given in the previous example, and suppose that over the entire time horizon, 4000 of system A were produced, 2000 of system B were produced, and 1000 of system C were produced. If PF1 has a startup cost of $700 in time period 1, then that cost will get allocated in that time period as follows.

	Component 1		Component 2	
	M1	**M2**	**M1**	**M2**
A	$200	$100	$20	$80
B	$20	$20	$80	$80
C	$0	$52.63	$0	$47.37

7. Future program expenditures (i.e., *startup, per-period active, and RDT&E costs*) are allocated to mission roles based on the proportion of total systems fielded from the future program (across the entire time horizon) that go to that mission role.

6. SCHEDULING CONSTRAINTS

Analysts can specify several scheduling constraints that influence the fleet transformation schedule.

6.1. Minimum Modernization Schedules

The analyst may want to require that certain systems be removed from the fleet by a particular time. The minimum modernization schedule allows the analyst to specify what percentage of systems in the original fleet must be modernized (upgraded or replaced) by certain times. A percentage can be supplied for each system type in each mission, for each time period. If at all possible, CPAT will modernize at least the specified percentage of systems on or before the indicated time period. CPAT may choose to modernize earlier than the schedule dictates, or in larger quantities, but will not modernize later or fewer than specified unless it is impossible to satisfy the modernization schedule. In fact, CPAT will violate budgets if necessary to satisfy the modernization schedule. This is so that analysts can always assess feasible courses of action (COAs) being considered by stakeholders, even if they are over budget (in which case the cost overages are captured and discussed as part of the COA results).

Note that minimum modernization schedules only apply to systems in the initial fleet.

6.2. Group Upgrade Limits

It may not be desirable to upgrade all groups simultaneously. For each mission across all components, the analyst can specify the maximum number of groups that can be upgraded per time period. Analysts can also specify the maximum number of groups that can be upgraded in total throughout the study horizon.

6.3. Upgrade Density

It may not make sense to utilize too few of a particular system. For each mission across all components, the analyst can specify up to three *minimum upgrade densities*, which define allowable numbers of groups that must use the same type of system for that mission at some point in the study, if it is used at all. The groups do not have to use that same system type concurrently, they just have to upgrade to the same system type at some point.

If density level(s) are specified, then the number of groups which use the specified system type at some point must be a) zero, b) *exactly equal* to one of the lower densities, or c) *at least as great* as the largest specified density. In other words, the number of groups can exceed the largest density, but otherwise must be exactly zero or one of the specified densities.

> Example: If mission M has a level-1 density of 3, then if any group upgrades mission M to system S in any component, then at least 2 other groups across all components must also upgrade mission M to system S. There are no constraints on when those additional upgrades must take place, and they may even occur after the first group has already moved on to yet another system type. The upgrade density constraint would also be satisfied if no group ever upgraded to system S.

Example: If mission M as a level-1 density of 3 and a level-2 density of 7, then if any group upgrades mission M to system S in any component, then either exactly 2 OR at least 6 other groups must also upgrade mission M to use system S in across all components.

Example: If mission M as a level-1 density of 3, a level-2 density of 7, and a level-3 density of 10, then if any group upgrades mission M to system S in any component, then either exactly 2 OR exactly 6 OR at least 9 other groups must also upgrade mission M to use system S across all components.

6.4. Final Density

The analyst can specify up to three *final system densities* for each mission. This represents the allowable numbers of groups that must be using any system type present in the final fleet. Similar to Upgrade Density, the final density constraint may be satisfied by either having zero groups using a given system, by having exactly the number of groups specified by one of the final densities, or by having more groups using that system than the largest specified final density.

Note that minimum final density constraints do not apply to systems in the initial fleet, only to systems that are the result of purchases/upgrades. If initial systems are still present in the final fleet, they do not have to be used by any minimum number of groups.

Example: Assume mission M has a minimum final density of 4. If any group uses system S for mission M in any component in the final fleet and S was not a part of the initial fleet, then at least 3 additional groups must also use system S for mission M across all components in the final fleet. Note that zero groups using system S would also satisfy this constraint.

Example: Assume mission M has two final densities specified, 2 and 4. If system S was not part of the initial fleet, then mission M may have 0, 2, 4, or more than 4 groups using system S across all components in the final fleet. There may not be 1 or 3 groups using system S in the final fleet.

Example: Assume mission M has three final densities specified 2, 4, and 7. If system S was not part of the initial fleet, then mission M may have 0, 2, 4, 7, or more than 7 groups using system S across all components in the final fleet. There may not be 1, 3, 5, or 6 groups using system S in the final fleet.

6.5. Final Population

The analyst can force a specific system to be present in the final fleet. For each mission in each component, the analyst can specify the *minimum final count* for each system type. CPAT will select upgrades that cause at least the specified number of systems to be in service to that mission in that component in the final time period.

6.6. Synchronization Sets

There are situations when upgrades across multiple missions in a single component must occur simultaneously. CPAT supports this through *synchronization sets*. A synchronization set consists of a set of missions in a component and a set of system types. The number of groups using systems from the synchronization set must be the same for all missions in the synchronization set component in the synchronization set at all times. If the number of groups using systems from the synchronization set changes for one mission in the synchronization set, then the same number of groups must make similar changes for all other missions in the synchronization set. This causes changes to be synchronized across all missions in the synchronization set. The missions do not necessarily need to use the same system types; they just need to all be using systems from the synchronization set.

At the beginning of the study horizon, the number of groups using systems in the synchronization set must be the same for all missions in the set.

Example: A synchronization set for Component 1 includes missions M1, M2, and M3, and systems S1, and S2. None of the missions is using S1 or S2 at the beginning of the study. In time period 4, 3 groups begin using S1 for mission M1 in Component 1. To stay in sync, an equal number of groups (3) must begin using either S1 or S2 for M2 and M3 in Component 1 in that same time period.

Later, in time period 6, M1 stops using S1 in Component 1 and starts using S2 instead. No changes are required in the other missions because both S1 and S2 are in the synchronization set; the number of groups using synchronization set systems has not changed.

In time period 10, two groups stop using S2 for M1 in Component 1 and start using S4, which is not in the synchronization set. Two groups must also switch to non-synchronization set systems for M2 and M3 in Component 1.

Although this example was written as if M1 in Component 1 were driving changes and M2 and M3 also in Component 1 were following M1's lead, this is not strictly the way it works. The changes to all missions in the component do have to occur simultaneously, but no particular mission is driving the others. A change in *any* mission in the synchronization set is contingent upon compatible changes in all other synchronization set missions.

6.7. Component Mission Succession

Missions in any component can be designated to succeed one another mission in the same or different component. That is, a mission in a component can be selected to follow another mission in another component so that nothing can be fielded to the succeeding mission until the preceding mission has 1) completely finished fielding and 2) modernized 100% of its original systems.

Example: Assume that there are two missions, M1 and M2. M1 is in Component A and M2 is in Component B. Suppose that M1 in Component A is selected to precede M2 in Component B. If M1 in Component A does not modernize all of its initial systems or it never completely finishes fielding, then M2 in Component B is not allowed to modernize any of its original systems. If M1 in Component A completes fielding and modernizes 100% of its systems in time period 12 then the earliest that M2 in Component B could modernize any of its systems is also time period 12.

6.8. Economic Useful Life

All systems in the fleet are given an economic useful life. This is the maximum age the system is allowed to be before it must be retired from a mission. Since transitions are made at the group level, all the systems in a group are assigned the same "average" age. Once a group of systems exceed their economic useful life then they must be retired from service. All initial groups in service with the same set of system types are given the same initial average age.

Example: Assume that System A has an economic useful life of 20 years. Two groups of System A are initially in Mission 1 and both groups have an average age of 10 years at the beginning of the time horizon. By time period 10 both groups of System A in Mission 1 must be retired.

Purchased or upgraded systems are given an age of zero. Systems do not start aging until they are assigned to a mission.

Example: Assume that System A has an economic useful life of 20 years. Two groups of System A are purchased over the course of time periods 1 and 2 and put into storage. In time period 4, these groups are fielded to Mission 1. In addition, one group in Mission 1 is also upgraded to System A in time period 4. All three of these groups of System A modernized in time period 4 in Mission 1 must be retired from Mission 1 before time period 24.

7. PRIORITY TIERS

There are situations when not all missions should be given equal consideration. For example, one set of missions might represent a high priority (e.g., deployable brigades) that needs to have the greatest possible capabilities, while a second set might represent lower priority resources that should only be upgraded if they do not interfere with high priority improvements. CPAT supports this scenario through priority tiers.

Each mission is assigned a priority, from 1 to N. Missions with priority 1 have the highest priority, while missions with priority 2 have the next highest priority, and so on. Multiple missions can have the same priority, and unless the analyst changes it, all missions have a priority of 1. The collection of missions with the same priority are said to be in the same *priority tier*. For example, all missions with priority 1 are in priority tier 1, or simply tier 1.

When identifying an ideal modernization schedule, CPAT considers each tier one by one. First it finds the optimal modernization schedule for the highest priority tier while preventing any changes to the fleet in lower tiers. It then optimizes the next highest tier using any leftover budget and production resources, finding an optimal modernization schedule for that tier while honoring the modernization schedule for the higher tier, and holding the fleet composition constant for lower tiers. For each tier, the modernization decisions made for earlier tiers are respected, and modernization of later tiers is temporarily prevented.

All business rules must be honored during the optimization of each tier. Constraints applying to production gaps, capacities, and budgets apply to each tier, taken in conjunction with decisions made for earlier tiers. For example, the highest tier's modernization schedule must be able to be implemented without causing production gaps for any product family. The next tier's schedule must also avoid production gaps, which means that its production must be just before, just after, or concurrent with the upper tier production schedule.

It is worth noting that the multi-tier approach may not give the globally optimal solution for the fleet overall, and may not even give the "true" optimal solution for an upper tier as it cannot take advantage of cost and production synergies with later tiers. This was a deliberate design decision: each tier's modernization schedule must be implementable regardless of what a lower tier does, even if that results in lower total performance. Due to the nature of what the tiers represent, it was important that plans for a given tier not rely on any lower tier. Each tier does, however, rely on upper tiers' schedules, and may need to make adjustments if plans in an upper tier are changed.

8. CONSTRAINT VIOLATIONS AND INTRA-TIER OPTIMIZATION PHASES

CPAT attempts to find production and fielding schedules which satisfy all business rules. It may be impossible to generate a schedule which satisfies all constraints. Sometimes there may be more than one way to resolve a conflict, such as a choice between going over budget and falling behind schedule. Furthermore, there may be multiple solutions with no constraint violations and equal fleet performance.

CPAT attempts to address these conflicts and ambiguities by utilizing six separate optimization phases. The user can determine which phases are executed, and in which order, depending on the analysis question being answered. The optimized value in each phase becomes a constraint on later phases. The phases are:

- Minimize schedule mandate violations
- Minimize economic useful life violations
- Minimize yearly budget violations
- Minimize horizon (cumulative) budget violations
- Maximize fleet performance
- Minimize cumulative combined costs (a user-chosen combination of cumulative procurement, O&S, and RDT&E expenditures)

If the user selected all six phases in the order above (not typical), they would be executed as follows. The first would minimize schedule violations; the second would minimize economic useful life violations while not allowing schedule violations to increase, the third would minimize yearly budget violations while not allowing either schedule or economic useful life violations to increase; the fourth would minimize horizon budget violations while not allowing schedule, economic useful life, or yearly budget violations to increase; the fifth would maximize fleet performance while not allowing schedule, economic useful life, yearly budget, or horizon budget violations to increase; the sixth would minimize cumulative combined fleet costs while preserving fleet performance and not allowing schedule, economic useful life, yearly budget, or horizon budget violations to increase.

Mandates, economic useful life constraints, and budget violations were all chosen to be minimized penalties rather than hard constraints. This ensures that if such business rules must be broken, the phase ordering prevents tradeoffs between violations and the user can diagnose the issue. Costs are often minimized after the performance maximization phase to ensure performance is achieved via the most intelligent possible allocation of budget resources and lower tiers, which use the left-over budget from higher tiers, will have the best possible opportunity for modernization.

These phases are applied within a single tier. In a multi-tier model, once a tier's schedule has been identified, the tier's fielding decisions cannot be changed by later tiers, even if doing so would resolve a constraint violation for the later tier.

CPAT also provides analysts with the ability to enforce hard constraints on any of the phase objective metrics. For example, the analyst may specify that the total number of economic useful life violations must be less than X, or that the fleet's cumulative performance must be greater than Y. These limits must be satisfied in all optimization phases. This capability allows the analyst to answer a myriad of questions such as “what is the cheapest fleet modernization plan with less than X economic useful life violations?” and “what is the smallest horizon budget violation required to achieve a fleet with at least X cumulative performance?”

9. FUTURE SYSTEMS AND PROGRAMS

The fleet structure also includes *future systems* which are similar to conventional systems with the following simplifying assumptions:

- The analyst defines a delivery schedule for each future system.
- Each future system is assigned to exactly one mission.
- Once a future system is delivered then no other non-future system transitions are allowed to the mission in the component where the future system was delivered.

The analyst specifies the transitions between systems and future systems along with the per-unit procurement costs, administrative delay, production delay, and long lead cost ratio. In addition, future systems have a performance rating determined by their mission, future systems have an O&S cost, and future systems may have an LRIP schedule. A future system may optionally be mandated so that it is forced to field.

Future systems can be assigned to any number of groups called *future programs.* Future programs are similar to product families and contain collections of future systems with similar characteristics. The analyst can specify whether future programs have start-up costs, an RDT&E Effort (this includes any per-period RDT&E cost, and is not allowed to be delayed), or a per-period procurement active cost. Future programs can optionally be designated by the analyst as "all or nothing," which specifies that if the program is activated, then every future system assigned to the program must be fielded.

Note that future systems and future programs still adhere to all business rules but their behavior is restricted due to the assumptions made above. By convention, a mission will never be supported by both future systems and systems that are allowed to coast.

10. POST-PROCESSING CALCULATION OF AVERAGE AGE

Several CPAT results require the calculation of the average age of all systems of a given type that are serving a certain mission in a certain component during a specified time period. Note that this calculation is done during post-processing – after the optimization has finished. This is distinct from the calculation of Economic Useful Life violations, which is needed during the Age Phase of the optimization. The average age calculation is based on the following assumptions:

1. System types that are initially in the fleet are, by convention, not able to be procured via purchases or upgrades. If System A is in the initial fleet and is also available for procurement, then this procured system is modelled as a different system type (e.g., System A-New).
2. Systems do not age while they are in storage, and retired systems do not carry their age with them into storage. Therefore, whenever any system is fielded, it is assumed to be brand new. Potential issues that could result from this assumption are partially mitigated by the Fielding New Systems from Storage business rule. Additionally, retired systems typically undergo some level of refurbishment before they are re-fielded, so this assumption is not unrealistic.
3. When systems are retired from a mission in a component, the oldest systems of that type supporting the given mission in the given component are assumed to be retired first.

For a given time period, the average age of the systems of a certain type serving a given mission in a given component is calculated by taking the sum of the ages of the appropriate systems divided by the total number of those systems.

Example: Consider three systems, S1, S2, and S3 in a fleet with only one component, C1. System S1 supports Missions M1 and M2, while S2 and S3 support Mission M3. In the initial fleet, there are 60 S1 with an initial age of 3 supporting M1, 30 S1 with an initial age of 3 supporting M2, and 100 S2 with an initial age of 5 supporting M3.

20 units of S1 are retired from M1 in each of the first three time periods. In time periods 4 and 5 respectively, 10 and 20 of these retired S1 systems are re-fielded to Mission M2. Additionally, 100 units of System S3 are purchased in time period 1. These systems are fielded to Mission M3, replacing 20, 30, 40, and 10 units of S2 in time periods 3, 4, 5, and 6, respectively. 30 units of S3 are retired from M3 in time period 7.

	Initial	TP1	TP2	TP3	TP4	TP5	TP6	TP7
S1/M1/C1	60	40	20	0	0	0	0	0
S1/M2/C1	30	30	30	30	40	60	60	60
S2/M3/C1	100	100	100	80	50	10	0	0
S3/M3/C1	0	0	0	20	50	90	100	70

Number of Systems in the Fleet

	Initial	TP1	TP2	TP3	TP4	TP5	TP6	TP7
S1/M1/C1	3	4	5	0	0	0	0	0
S1/M2/C1	3	4	5	6	5.5	4.67	5.67	6.67
S2/M3/C1	5	6	7	8	9	10	0	0
S3/M3/C1	0	0	0	1	1.40	1.78	2.60	3.14

Average Age of Systems in the Fleet

11. BUSINESS RULES

This is a consolidated list of the business rules incorporated into the CPAT optimization model. Most of these business rules are described in context in the documentation above.

11.1. System Transition Flow

11.1.1. Constant System Population

Throughout the planning horizon, each mission in each component always maintains a constant number of systems. Every change to the fleet consists of either modifying existing systems or removing some number of systems from the fleet and putting an equal number of different systems in their place.

11.1.2. Group Purity

At any given time, the systems serving a particular group for a particular mission in a component must all be of the same system type. Different groups in the same or different components can each be using a different system type for that mission, and different missions within the same group may be using different system types, but a single group cannot mix system types within the same mission.

11.1.3. Upgrade and Exchange Availability

For any time period, the number of systems of a given type in a mission in a component that are upgraded or sent to storage may not exceed the number currently *modifiable*. Similarly, the number of systems of a given type in storage that are upgraded or sent to a mission in a component may not exceed the number currently *modifiable*. In both cases, the number of currently *modifiable* systems is given by the current number present minus the current number that is in the process of being upgraded (which may take several time periods due to administrative and production delays).

11.1.4. Initial Populations

Each mission in each component has an initial population of systems that is already in the fleet and is immediately available to begin modernization. There may also be an initial population of systems in storage which are assigned to components also immediately available to begin upgrading or swapping into missions in that component.

11.1.5. Storage Flow

Systems enter and exit storage through the following means: 1) purchases put new systems directly into a particular component in storage, 2) storage upgrades take one system type already in storage and turn it into another type in the same component, and 3) storage swaps take one system type out of a mission in a component and into storage into the same component while taking another type out of storage and sending it to the mission in that component. Once in storage, a system is immediately available for any type of flow action with one exception: a system cannot be swapped into and out of the same mission in the same component in the same time period.

11.1.6. No Pre-Usage Upgrades

Newly purchased systems in storage in any component that have not yet been sent to a mission should not be upgraded.

11.1.7. Optional Pre-Purchasing

Systems may be purchased or in-storage upgraded before they are actually needed to be fielded to a mission in a component. However, this ability is optional and may be disallowed by user choice.

11.1.8. Optional Storage Upgrades

Systems may be upgraded in mission or in storage for any component. The user has the option to specify for each upgrade whether that upgrade can be done in storage.

11.1.9. Delivery Implies Fielding

System types whose procurement cost is non-zero can only be delivered to a component if they are also subsequently fielded to a mission in that component. (Note that delivery of these systems from production and fielding can occur at different times.) Only system types that can be procured for free (usually hull systems) can be delivered to a component without also being fielded.

11.1.10. No Retire and Re-Fielding

Systems that are retired from a mission and sent to storage cannot be immediately sent back into that same mission during that same time period.

11.1.11. One-Year Duty Minimum

Systems in a mission in any component must remain for at least one time period before they can be swapped out or spoken for by a mission upgrade.

11.1.12. Fielding New Systems from Storage

New systems delivered to storage via purchases or in storage upgrades may only remain in storage for a user defined number of time periods. Additionally, any LRIP systems or systems in initial storage may only remain in storage for the user-specified number of time periods, beginning in the first time period that non-LRIP production is completed.

11.1.13. Disallow Instantaneous Cross-Mission Transfers

Systems retired from a mission in a component in any time period cannot immediately be re-fielded to another mission in a component in the same time period.

11.2. Mission Priority Tiers

11.2.1. Priority Tiers

Fleet missions in any component may be partitioned into priority tiers across components wherein each tier comprises a separate optimization. The modernization of missions in all components in the highest priority tier is performed first, with subsequent tiers being modernized separately with the remaining budget. Note that all other business rules must hold *in toto* across all tiers. For example, if a product family disallows production gaps, then it may only be started up once even if it fields systems to missions in components across multiple tiers; it is not allowed to start up separately for each tier.

11.2.2. Tier Phases

Within each tier, there are six separate optimization phases. These phases are described in detail in Section 8.

11.2.3. Hard Limits on Tier Phases

CPAT provides analysts with the ability to enforce hard constraints on any of the phase objective metrics. Examples of this capability can be found in Section 8.

11.2.4. Component Mission Succession

One mission in a specific component can be designated to succeed another mission in the same or a different component so that nothing can be fielded to the succeeding mission in that component until the preceding mission in the specified component has 1) completely finished fielding and 2) modernized 100% of its original systems. This is typically used for corresponding missions in different components (i.e., the IFV mission in the Active Army component must be fully upgraded before the IFV mission in the National Guard component), but can also be used for missions in the same components.

11.3. Transition Delays

11.3.1. Delay Partitioning

When upgrading from one system to another (whether in a mission or in storage for any component) or purchasing a new system, there may be a delay between when the new system is paid for and when it is delivered. This delay is partitioned into an administrative delay (where the system has been paid for but is not yet in production) followed by a production delay (where the system is in production but is not yet delivered). These delays must be accounted for. Default administrative and production delay = 0 periods.

11.3.2. Upgrade Administrative Delays

For any upgrade having administrative and production delays, the administrative period is allowed to begin even if the seed system is not yet on hand. However, the seed system <u>must</u> be on hand to begin the first production period. Intuitively, this means that “upgrade paperwork” (i.e., the administrative period) can be started in anticipation of the soon to arrive system. Stated another way, while in administrative periods a system is not yet “spoken for.”

11.4. General Scheduling Rules

11.4.1. System Modernization Requirements

Some system types in the initial fleet require that a certain percentage must be transitioned to some other system type at or before a specified time in the planning horizon. This percentage is applied across the fleet and is not specific to systems in each component. This modernization must be performed. Default requirement = 0%.

11.4.2. System Mandates

For some missions in some components, a minimum number of a particular system type is mandated to be in that mission in that component by the end of the planning horizon. This minimum must be met. Default minimum = 0.

11.4.3. Per-Period Mission Modernization Limit

For certain missions, an upper bound may exist on the number of groups that are allowed to modernize that mission across all components in a single time period. These upper bounds must be respected. Default bound = unlimited.

11.4.4. Cumulative Mission Modernization Limit

For certain missions, an upper bound may exist on the cumulative number of groups that are allowed to modernize that mission across all components throughout the entire planning horizon. These upper bounds must be respected. Default bound = unlimited.

11.4.5. Minimum Group Transition Density

For each mission, if a system within that mission transitions to another system, there may be up to 3 density levels that dictate how many groups must be transitioned in this manner across all components.

- *Example 1:* Levels = {12,-,-} implies group transition density must be at least 12.
- *Example 2:* Levels = {12, 16,-} implies group transition density must be either exactly 12 or at least 16.
- *Example 3:* Levels = {12, 16, 20} implies group transition density must be exactly 12, exactly 16, or at least 20.
- Default Levels = {-,-,-}.

11.4.6. Minimum Group Final Density

Missions may require that the number of groups of non-initial systems in the mission during the final time period across all components meet certain densities. These densities may be specified by up to 3 levels, which operate analogously to the Minimum Group Transition Densities. Default Levels = {-,-,-}.

11.4.7. System Obviation

Each system type may be obviated by any number of other system types. A system may only be delivered earlier than the earliest delivery of any of its obviating system types, with the option to specify a number of overlap time periods in which both systems can be delivered. This applies to systems in any mission and in any component

11.4.8. Synchronization Sets

A collection of missions in a single component may contain a collection of systems that must all modernize simultaneously. These systems and mission would be assigned to a synchronization set. If a certain number of groups of systems in one mission in one component are modernized, then the same number of groups of any synchronized systems must also modernize in any other synchronized missions in that same component in that same time period.

11.4.9. Storage Upgrade Consumption Priority

Certain systems in storage in any component may take upgrade consumption priority over certain other systems. This means that for each component in storage if the higher priority system in that component is exchangeable in storage, then, it must be used as an upgrade seed before the lower priority system in that component can be used as an upgrade seed.

11.4.10. Upgrades Trump Purchases

For some systems in each component, modernization must be accomplished via upgrades, if possible. For each component, a new purchase is allowed only if no seeds systems are available for the upgrade in that component.

11.4.11. Economic Useful Life

Certain systems may have an upper bound on the number of time periods that they can spend in a mission in any component. Once a system reaches that specific limit or age then it must be retired from service. This rule does not apply to terminal systems (i.e., systems that cannot be transitioned to any other system).

11.5. Budgets

11.5.1. Per-Period Budgets

The amount of money spent each period in the 3 categories of Procurement, O&S, and RDT&E must not violate associated yearly budgets for these expense types. Furthermore, a user-specified combination of these 3 yearly budget types must not violate a yearly combined budget. These budgets must be respected by both future and non-future system expenditures throughout the conventional and extended time horizons. Default budgets = unlimited.

11.5.2. Cumulative Budgets

The total amount of money spent throughout the planning horizon in the 3 categories of Procurement, O&S, and RDT&E must not violate associated cumulative budgets for these expense types. Furthermore, a user-specified combination of these 3 budget types (matching the per-period budget combination) must not violate a combined cumulative budget. These budgets must be respected by both future and non-future system expenditures. Default budgets = unlimited.

11.5.3. Component Earmarks

Additional money, above the per-period budget, may be allocated to purchases and upgrades of systems in a specific component in a specific conventional or extended time period. An earmark for a component does not require purchases and upgrades of at least that amount to be spent for that component, but it does not allow a portion of the earmark to be spent for the purchases and upgrades in another component. Default earmarks = 0.

11.5.4. Product Family Earmarks

Additional money, above the per-period budget, may be allocated to active costs, startup costs, LRIP, RDT&E, purchases, and upgrades of systems in a specific product family in a specific conventional or extended time period. An earmark does not require at least that amount to be spent for that product family, but it does not allow any portion of the earmark to be spent for active costs, startup costs, LRIP, or RDT&E of any other product family or for the purchases and upgrades of systems not in the specified product family. When a system is purchased or upgraded in a time period in which it could charge either a product family earmark or a component earmark, the product family earmark is charged. Default earmarks = 0.

11.5.5. Early/Late Transition Charging

No transition may take place in a time period early enough so that associated costs (whether transition, long lead, or product family start-up costs) would be incurred prior to the start of the time horizon. Similarly, no transition may occur in time periods late enough that associate product family start-up costs would be incurred after the end of the time horizon.

11.5.6. Long Lead

Some system types may have long lead on their procurement. This means that a certain percentage (long lead cost ratio) of their procurement cost is incurred one year earlier than normal. (Remember that normally procurement costs are incurred during the first administrative period.)

11.6. Product Families

11.6.1. Active Product Families

Multiple system types can be clustered together into a single product family, with the interpretation that these systems share production facilities. A product family is considered “active” (thus incurring per-period procurement and RDT&E costs) during a time period if any member systems are 1) in administrative delay, 2) in production delay, or 3) being delivered and the production delay is 0. Note that <u>both LRIP and FRP</u> count towards these three conditions, even if the LRIP is being incurred for a separate product family.

11.6.2. Family Start-Up Costs

Each product family may have an associated start-up cost profile that must be incurred when the family first begins work for full-rate production. That is, when the family is 1) in administrative delay, 2) in production delay, or 3) being delivered and the production delay is 0 for the first non-LRIP systems. These costs are allocated to missions in components using a system density weighting method. Default start-up cost = $0.

11.6.3. Family Per-Period Costs

Each product family may have an associated per-period procurement cost and/or per-period RDT&E cost that must be incurred every time period that the family is active. Note that a family is active even if its member systems are being produced for LRIP of another family. Like Family Start-Up Costs, these costs are allocated to missions in components using a system density weighting method. Default per-period cost = $0.

11.6.4. Family Minimum Per-Period Delivery

For each product family and time period, there may be a lower limit on the number of member systems delivered during that period. These limits must be respected, although LRIP does not count towards this capacity. Default capacity = 0.

11.6.5. Family Maximum Per-Period Delivery

For each product family and time period, there may be an upper limit on the number of member systems delivered during that period. These limits must be respected, although LRIP does not count towards this capacity. Default capacity = unlimited.

11.6.6. Family Minimum Cumulative Capacity

For each product family, there may be a lower limit on the cumulative number of member systems that are ever delivered from the family if any systems are delivered from the family. These limits must be respected. All produced LRIP and coasting systems count towards this capacity. Default capacity = 0.

11.6.7. Family Maximum Cumulative Capacity

For each product family, there may be an upper limit on the cumulative number of member systems that are ever delivered from the family. These limits must be respected. All produced LRIP and coasting systems count towards this capacity. Default capacity = unlimited.

11.6.8. Minimum Sustaining Rate

Given that systems are delivered from a product family in a particular time period, there may be a lower bound on the number of systems that must be delivered from that family in that time period. These bounds must be met, although LRIP does not count towards this bound. Also, these bounds are not enforced during the last production period, allowing the production line to wind down. Default MSR = 0.

11.6.9. Delivery Gaps

Product families may be restricted so that delivery begins only once; it cannot start delivering systems, stop, and then subsequently restart. This means that all systems within that family must be delivered during a collection of contiguous time periods.

11.6.10. Production Smoothing

For each product family, there may be a limit on the variation in the number of systems delivered from that family when in full-rate production. This prevents undesirable effects to the manufacturer. Note that in the final period of full-rate production, this restriction is not enforced so that the production line can begin to wind down output. Default production variation = unlimited.

11.6.11. Production Ramp-Up

For each product family, there may be a ramp up period prior to full-rate production. During this ramp-up, delivery output is not required to respect production smoothing. Instead, the number of systems delivered must be non-decreasing in time during this ramp-up.

11.6.12. Upfront RDT&E Cost

For each product family, there may be an upfront RDT&E cost such that systems from the family can be delivered if and only if the RDT&E cost profile of the family is incurred. Default cost = \$0. The analyst may choose to allow the optimization engine to delay certain upfront RDT&E costs to avoid budgetary bottlenecks. For each time period that a cost profile is delayed, a separate cost profile must be supplied; a delay d (including d=0) is valid only if it has an associated cost profile. Incurring a delay of d time periods also delays the availability of systems in the product family by d time periods. In addition, if d>0, then at least one system within that family must also be delayed by exactly d (other systems may be delayed by more). The analyst may choose to enable legacy RDT&E cost behavior. As before, systems from the product family with an upfront RDT&E cost profile may be produced if and only if the cost profile is incurred. However, the d=0 cost profile is incurred regardless of when the associated systems are first delivered. An upfront RDT&E cost profile cannot be incurred if any of the costs extend into future time periods.

11.6.13. Product Family Obviation

All systems within a product family may be obviated by any system within another product family so that any system from the obviated product family can only be delivered prior to any deliveries of systems from the obviating product family.

11.6.14. Product Family Ratios

For each product family, there may be a ratio defined for each component in the model in which delivery of systems from the product family to each component must be within a set variance of the defined ratios. The analyst also specifies a time period window in which the delivery ratio must be met. Once a product family starts delivering systems to one component, it must also begin delivering systems to all other components according to the ratio. Once delivery is complete to one component, the ratio for that component is no longer enforced. However, no more systems from this product family may be delivered to that component after delivery has ceased.

11.7. Low-Rate Initial Production

11.7.1. LRIP Profiles

Some systems in some product families may require a modest number of systems be produced in the years leading up to full rate production for the family. These LRIP profiles define fixed numbers of systems that must be produced up to 5 years before FRP begins. These LRIP profiles have 3 additional analyst-defined properties: 1) Not all of the LRIP systems produced have to be delivered to storage (some may be destroyed, for instance), 2) the seed system for the LRIP may or may not be explicitly defined and, 3) if the seed system is defined, these seeds may or may not be extracted from storage when the LRIP profile is produced.

11.7.2. LRIP Timing

All LRIP profiles incurred by a product family must be lined up so that their final LRIP delivery occurs exactly one time period prior to the first non-LRIP (i.e., FRP) delivery for the family.

11.8. Coasting Systems

11.8.1. Coasting System Fielding

Some systems, designated by the user, are allowed to continue delivery to missions in any component in the extended time horizon based on their delivery rates in the last conventional time period. These systems are not considered future systems and have more flexibility as their delivery rates are selected by the optimization. If the optimization chooses to coast a system in a mission in a component, then it coasts that system until the mission has no more groups in which to upgrade to the coasting system. Once delivery of coasting system to a mission in a component ceases, then it cannot be restarted in a later time period. Coasting systems are purchased or upgraded based on how they were acquired in the last conventional time period. By convention, a mission will never be supported by both future systems and systems that are allowed to coast.

11.9. Future Programs

11.9.1. Future Program Activation

Systems that might enter the fleet far in the future can be grouped together into future programs. Future programs are incorporated into the fleet via simple go/no-go decisions. If a future program is activated, then at least one future system associated with the program must be activated. Optionally, each future program may be restricted so that its activation requires that all of its associated future systems be fielded.

11.9.2. Future System Fielding

When a future system is activated in a component, it must be fielded to its mission in that component according to a fixed, user-defined fielding schedule. Optionally, each future system may be mandated to be fielded, in every component. If mandated future systems do not field then the schedule phase indicates infeasibilities.

11.9.3. Future Obviates Present

Once a future system starts fielding to a mission in a particular component, no other "non-future" systems may be fielded to that mission in that same component.

GLOSSARY

Class	A partition of the systems or missions in the fleet that provides a useful aggregation of certain optimization results for ease of comprehension. Class partitions do not have to follow set and group organization of the fleet and can be defined in any manner that is useful to the analyst. In essence, they provide “post-processing” organization of fleet results.
Component	A division of the fleet into separate operational units with similar structure.
Cost Profile	A set of costs scheduled to occur across multiple time periods. Times within a cost profile may either be relative to some event (such as a startup time period) or absolute (referring to specific time periods within the study horizon).
Fleet	A collection of systems, organized into components, missions, groups, and sets.
Group	A cluster of resources working together and supporting one or more missions.
In-Mission Upgrade	A transition which represents the modification of an existing system type within a mission into another system with different performance attributes. Here, the old system is consumed to create the new system.
In-Storage Upgrade	To modify a system that has been sent to the storage yard before redeploying it into another mission. An in-storage upgrade may only take place immediately before the resultant system is deployed via reapplication.
Mission	An operational responsibility of a group, which requires a fixed number of systems to fulfill.
Procurement	A term referring to expenses incurred in the process of modernizing systems. Upgrade, purchase, reapplication, product family start-up and per-period procurement costs all fall under this category.
Product Family	A set of system types that share production costs, RDT&E costs, and/or resources.
Purchase/Replacement	A transition which does not consume the old system to produce the new system. The old system is sent to the storage yard and is available to be repurposed.
RDT&E	Research and other upfront activities that must take place if any associated systems are to be fielded. Per-period RDT&E costs also fall under this category. The acronym stands for Research, Development, Testing, and Evaluation.

RDT&E Effort	A cost profile and per-period costs associated with a collection of system types, representing the cost of upfront and per-period RDT&E activities needed to make the related systems available. If any system associated with an RDT&E Effort is ever fielded, the RDT&E Effort's cost profile is incurred. If any system associated with an RDT&E Effort is active, the RDT&E Effort's per-period cost is incurred for that period.
Reapplication/Repurposement	A transition that represents the deployment of system from the storage yard into a mission. The deployed systems are removed from storage and placed into service, and the displaced systems are removed from service and sent to the storage yard. Reapplication is often immediately preceded by an in-storage upgrade.
Set	A collection of groups, each of which supports the same set of missions.
Storage Yard	A conceptual holding area for systems which have been removed from service and have not yet been repurposed.
System	A resource which may be applied to a mission. Systems are the individual components which are being considered for upgrading and replacement by the CPAT optimization algorithm.
System Transition	A general term referring to any substitution event wherein one system type in service to a mission is switched over to another type. This conversion may happen via an in-mission upgrade, a purchase, or a reapplication.

DISTRIBUTION

External

1 Roy E. Rice

Teledyne Brown Engineering

300 Sparkman Drive

Huntsville, AL 35805-1912

1 Frank M. Muldoon

Applied Materials

5225 West Wiley Post Way, Suite 275

Salt Lake City, UT 84116

Internal

1	MS 1188	Dean A. Jones, 8830 (electronic copy)
1	MS 1188	Bruce M. Thompson, 8833 (electronic copy)
1	MS 1188	Craig R. Lawton, 8834 (electronic copy)
1	MS 1188	Alan S. Nanco, 8836 (electronic copy)
1	MS 1188	Dennis J. Anderson, 8836 (electronic copy)
1	MS1188	Lucas A. Waddell, 8833 (electronic copy)
1	MS 1188	Darryl J. Melander, 9365 (electronic copy)
1	MS 1397	Peter B. Backlund, 5837 (electronic copy)
1	MS 1188	Stephen M. Henry, 8833 (electronic copy)
1	MS 1188	Matthew J. Hoffman, 8831 (electronic copy)
1	MS1188	April M. Zwerneman, 8834 (electronic copy)
1	MS 0899	Technical Library, 10756 (electronic copy)

This page intentionally left blank

Sandia National Laboratories