SAND2015-4009

SANDIA REPORT

SAND Number Pending
Unlimited Release
Printed May 2015

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

Darryl J. Melander, Stephen M. Henry, Matthew J. Hoffman,
Gio K. Kao, Craig R. Lawton, Frank M. Muldoon, Roy E. Rice, Liliana Shelton

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

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

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@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND # Pending

Unlimited Release

Printed May 2015

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

[1]Darryl J. Melander, [2]Stephen M. Henry, [2]Matthew J. [2]Hoffman, [3]Gio K. Kao, [2]Craig R. Lawton, [2]Frank M. Muldoon, [4]Roy E. Rice, [5]Liliana Shelton

[1]Software Systems R&D, Org 9525
[2]System Readiness and Sustainment Technologies, Org 6133
[3]Networked System Survivability & Assurance, Org 5229
[5]System Surety Engineering I, Org 421

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

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

Abstract

To help effectively plan the management and modernization of its large and diverse fleet of vehicles, the Program Executive Office Ground Combat Systems (PEO GCS) 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 page intentionally left blank

Table of Contents

Table of Figures

Acronyms

CPAT	Capability Portfolio Analysis Tool
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

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 (Muldoon, 2015)[1].

[1] Muldoon, Frank M., Hoffman, Matthew J., Henry, Stephen M., Shelton, Liliana, Kao, Gio K., Melander, Darryl J., Lawton, Craig R. The Capability Portfolio Analysis Tool (CPAT): A Mixed Integer Linear Programming Formulation for Fleet Modernization Analysis. Albuquerque, NM: Sandia National Laboratories, 2015.

2 Fleet Structure

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

- 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 to carry out that mission. Each system is assigned to one and only one mission within a group.

This configuration is shown in Figure 1 below.

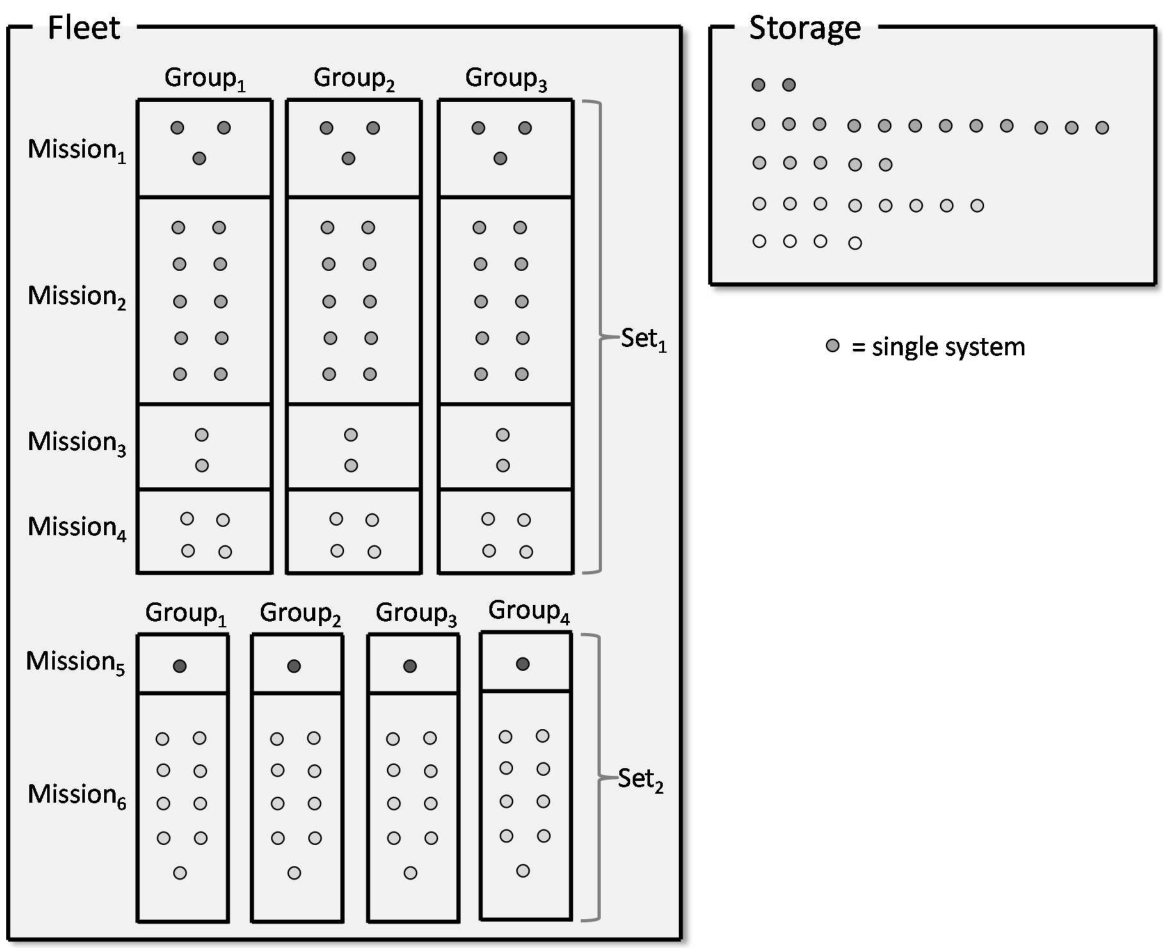

Figure 1: Fleet Structure

A fleet is described by defining each element of the fleet structure, including the 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 assigned to each mission at the beginning of the study horizon.

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. Each system type has a performance rating (composed of multiple performance categories) for each mission in which the system may serve. 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 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). 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 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 have differing system types across separate groups.

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, 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 Fraction – The portion of the procurement cost which is incurred one time period before the administrative period begins.

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.

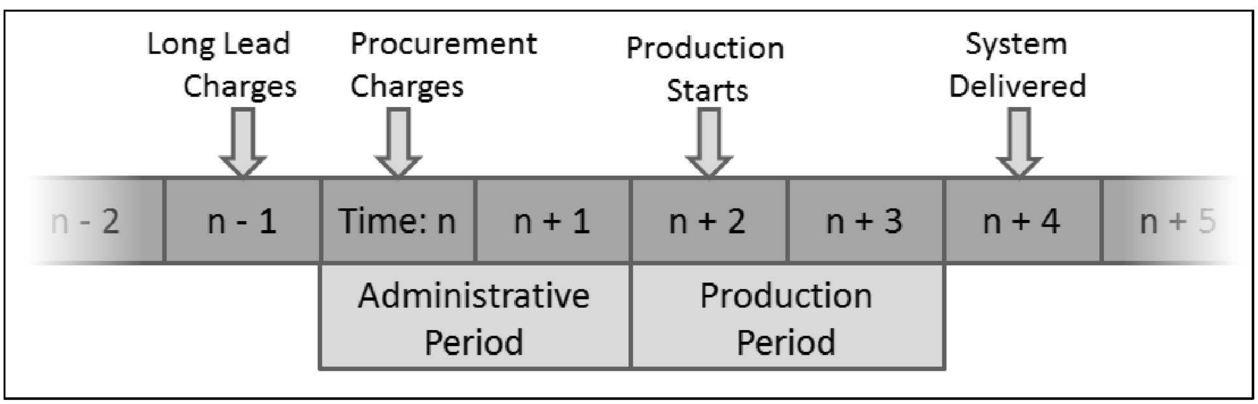

Figure 2: The Production Process

4.2 The Storage Yard

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

Furthermore, purchases can be made before the new systems are to be fielded to missions. Pre-purchased systems are also held 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.

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 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 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 Consumption Priority

It is possible for multiple different seed system types to be upgraded to the same "target" system type. 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 *consumption priority* among these seed systems. A *consumption priority* determines the sequence in which seed systems in storage are upgraded to any target system.

> Example: Assume that the optimization chooses to upgrade to System A. 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 from System B while there are System B's available in storage before upgrades from System C are made.

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, 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. 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 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 Capacity Schedules

The number of systems that can be delivered from a product family may be limited by production capacity or other real-world constraints. Each product family has a delivery capacity schedule. This is an analyst-specified, per-period limit on the maximum number of systems that can be delivered from the family each time period. The number of systems delivered from the family in a given time period may not exceed the specified delivery capacity.

Analysts may also specify a cumulative delivery limit for each product family. This represents the maximum total number of systems in the family that can be delivered throughout the entire study horizon.

Unless otherwise specified, delivery capacities are unlimited. As with delivery gaps, product family delivery limits take in-storage upgrades into consideration.

4.4.3 Product Family Minimum Delivery 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.

4.4.4 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 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.

4.4.5 Product Family Expenditures

Product families incur several types of expenditures: startup costs, per-period costs and RDT&E. The first two types are based on the product family's activity status; RDT&E is 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 is in an administrative or production period. A per-period expense is incurred each time period the family is active. A product family may also incur startup charges when it first becomes active. Startup charges are entered as a list of amounts to be charged, and the number of time periods before or after the 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.

4.4.6 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, while others will serve other purposes and can never be used by missions.

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 production expenditure 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.

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.

4.5 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.

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.

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 costs.
- Operating expenses represent the cost of using and maintaining systems present in the fleet.
- Research expenses represent the cost of any prerequisite activities that make new system types eligible for procurement.

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 systems are 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 fielding 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).

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

For reporting purposes, product family expenditures are allocated to missions using a vehicle density weighting method; the portion of each product family expenditure incurred by a given mission is proportional to the percentage of vehicles 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 vehicles.

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

5.2 Operating Expenditures

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 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 research before they become available. CPAT supports features which capture the costs of these prerequisite activities.

Each set of prerequisite research 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 costs are incurred. If no system associated with the family is ever in the fleet, the effort's costs are not incurred.

An RDT&E Effort's 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 cost profiles. Each profile is associated with a particular delay, and specifies the amounts and time periods when 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.

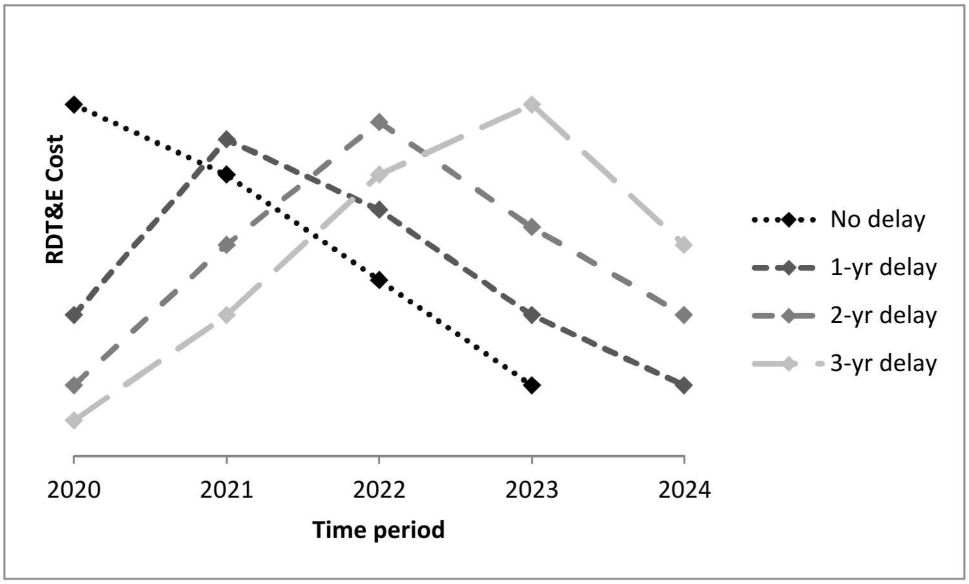

Figure 3: Notional RDT&E Cost Profile Example

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 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 RDTE costs would ever be charged.

RDT&E costs are distributed to missions using the same vehicle density weighting method as other product family costs.

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.

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, 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. Furthermore, there may be different "components" of the fleet, with different priorities, and it may be desirable to ensure that each component is either entirely modernized to a certain option, or not at all. For each mission, 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, then at least 2 other groups must also upgrade mission M to use 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, then either exactly 2 OR at least 6 other groups must also upgrade mission M to use system S.

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, then either exactly 2 OR exactly 6 OR at least 9 other groups must also upgrade mission M to use system S.

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 be 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 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 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 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 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, 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 the final time period.

6.6 Synchronization Sets

There are situations when upgrades across multiple missions must occur simultaneously. CPAT supports this through *synchronization sets*. A synchronization set consists of a set of missions 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 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 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. To stay in sync, an equal number of groups (3) must begin using either S1 or S2 for M2 and M3 in that same time period.

Later, in time period 6, M1 stops using S1 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 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.

Although this example was written as if M1 were driving changes and M2 and M3 were following M1's lead, this is not strictly the way it works. The changes to all missions 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 Mission Succession

Missions can be designated to succeed one another. That is, a mission can be selected to follow another mission 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. Suppose that M1 is selected to precede M2. If M1 does not modernize all of its initial systems or it never completely finishing fielding, then M2 is not allowed to modernize any of its original systems. If M1 completes fielding and modernizes 100% of its systems in time periods 12 then the earliest that M2 could modernize any of its systems is also time period 12.

7 Priority Tiers

There are situations when not all groups should be given equal consideration. For example, one group set might represent a high priority component (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 group set is assigned a priority, from 1 to N. A group set with priority 1 has the highest priority while priority 2 has the next priority and so on. Multiple group sets can have the same priority, and unless the analyst changes it, all group sets have a priority of 1. The collection of group sets with the same priority are said to be in the same *priority tier*. For example, all group sets 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 Priorities

CPAT attempts to find production and fielding schedules which satisfy all business rules. If it is impossible to generate a schedule which satisfies all constraints, CPAT will attempt to find a "least bad" schedule that violates constraints in the least impactful way. 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 applying the following priorities:

- *Prefer schedules with no violations over those with violations.* CPAT will only propose a schedule with constraint violations if it is unable to find a schedule without violations.
- *Prefer budget violations over schedule violations.* CPAT will make every attempt to satisfy modernization schedule requirements and minimum final count requirements. These schedule requirements will only be violated if the issue cannot be resolved by going over budget.
- *Prefer low cost schedules over high cost schedules.* If there are multiple solutions with the same performance and the same constraint violation characteristics, CPAT will prefer the schedule with the lowest combined cost.

These priorities are applied within a single tier. In a multi-tier model, each tier's schedule is identified using the priorities listed here. 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.

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 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 fraction. 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 is not allowed to be delayed), or a per-period active costs. 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.

10 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.

10.1 System Transition Flow

10.1.1 Constant System Population

Throughout the planning horizon, each mission 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.

10.1.2 Group Purity

At any given time, the systems serving a particular group for a particular mission must all be of the same system type. Different groups 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.

10.1.3 Outflow Availability

For any time period, the number of systems of a given type in a mission that are upgraded or swapped to storage may not exceed the number currently exchangeable. Similarly, the number of systems of a given type in storage that are upgraded or sent to a mission may not exceed the number currently exchangeable. In both cases, the number currently exchangeable is given by the current number present minus the current number in the process of being upgraded.

10.1.4 Initial Populations

Each mission 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 is also immediately available to begin upgrading or swapping into missions.

10.1.5 Storage Flow

Systems enter and exit storage through the following means: 1) purchases put new systems directly into storage, 2) storage upgrades take one system type already in storage and turn it into another type, and 3) storage swaps take one system type out of a mission and into storage while taking another type out of storage and sending it to the mission. 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 time period.

10.1.6 No Pre-Usage Upgrades

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

10.1.7 Optional Pre-Purchasing

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

10.1.8 Delivery Implies Fielding

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

10.1.9 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.

10.1.10 One-Year Duty Minimum

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

10.2 Mission Priority Tiers

10.2.1 Priority Tiers

Fleet missions may be partitioned into priority tiers wherein each tier comprises a separate optimization. The modernization of missions 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 across multiple tiers; it is not allowed to start up separately for each tier.

10.2.2 Tier Phases

Within each tier, there are four separate optimization phases. The first minimizes schedule violations; the second minimizes budget violations while not allowing schedule violations to increase; the third maximizes fleet performance while not allowing either schedule or budget violations to increase; the fourth minimizes cumulative combined fleet costs while preserving fleet performance and not allowing schedule or budget violations to increase. This ensures that 1) if business rules must be broken then budget violations are always preferred to schedule violations, 2) performance is achieved via the most intelligent possible allocation of budget resources and 3) lower tiers, which use the left-over budget from higher tiers, will have the best possible opportunity for modernization.

10.2.3 Mission Succession

One mission can be designated to succeed another 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.

10.3 Transition Delays

10.3.1 Delay Partitioning

When upgrading from one system to another (whether in a mission or in storage) 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.

10.3.2 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 must 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."

10.4 General Scheduling Rules

10.4.1 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 modernization must be performed. Default requirement = 0%.

10.4.2 System Mandates

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

10.4.3 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 in a single time period. These upper bounds must be respected. Default bound = unlimited.

10.4.4 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 throughout the entire planning horizon. These upper bounds must be respected. Default bound = unlimited.

10.4.5 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.

- *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 = {-,-,-}.

10.4.6 Minimum Group Final Density

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

10.4.7 System Obviation

Each system type may be obviated by any number of other system types. A system may only be fielded earlier than the earliest fielding of any of its obviating system types.

10.4.8 Synchronization Sets

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

10.4.9 Storage Consumption Priority

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

10.4.10 Upgrades Trump Purchases

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

10.5 Budgets

10.5.1 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. Default budgets = unlimited.

10.5.2 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. Default budgets = unlimited.

10.5.3 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.

10.5.4 Long Lead

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

10.6 Product Families

10.6.1 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 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 both LRIP and full-rate production (FRP) count towards these three conditions, even if the LRIP is being incurred for a separate product family.

10.6.2 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 using a vehicle density weighting method Default start-up cost = $0.

10.6.3 Family Per-Period Costs

Each product family may have an associated per-period 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 using a vehicle density weighting method. Default per-period cost = $0.

10.6.4 Family Per-Period Capacity

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.

10.6.5 Family 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, although LRIP does not count towards this capacity. Default capacity = unlimited.

10.6.6 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.

10.6.7 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.

10.6.8 Production Smoothing

For each product family, there may be a limit on the variation in number of system 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.

10.6.9 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.

10.6.10 RDT&E Cost

For each product family, there may be an RDT&E cost and 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 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 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.

10.7 Low-Rate Initial Production

10.7.1 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 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.

10.7.2 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.

10.8 Future Programs

10.8.1 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.

10.8.2 Future System Fielding

When a future system is activated, it must be fielded to its mission according to a fixed, user-defined fielding schedule. Optionally, each future system may be mandated to be fielded, in which case the schedule phase would be infeasible were the system not activated.

10.8.3 Future Obviates Present

Once a future system starts fielding to a mission, no other “non-future” systems may be fielded to that mission.

Glossary

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 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 costs all fall under this category.
Product Family	A set of system types that share production costs and 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 prerequisite activities that must take place if any associated systems are to be fielded. The acronym stands for Research, Development, Testing, and Evaluation.
RDT&E Effort	A cost profile associated with a collection of system types, representing the cost of prerequisite 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.
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

Internal

1	MS 0330	Liliana Shelton, 421 (electronic copy)
1	MS 0671	Gio K. Kao, 5629 (electronic copy)
1	MS 1188	Richard O. Griffith, 6130 (electronic copy)
1	MS 1188	Stephen M. Henry, 6133 (electronic copy)
1	MS 1188	Matthew J. Hoffman, 6133 (electronic copy)
1	MS 1188	Craig R. Lawton, 6133 (electronic copy)
1	MS 1188	Darryl J. Melander, 9525 (electronic copy)
1	MS 1188	Frank M. Muldoon, 6133 (electronic copy)
1	MS 1188	Alan Nanco, 6114 (electronic copy)
1	MS 1188	Bruce M. Thompson, 6133 (electronic copy)
1	MS 0899	Technical Library, 9536 (electronic copy)

This page intentionally left blank

Sandia National Laboratories