

Investment Analysis using the Portfolio Analysis Machine (PALMA¹) Tool

by Richard A. Moynihan
21 July 2005

Government Investment Analysis Guidance

Current Government acquisition guidelines mandate the development of strategic goals and related capability-based performance measures in making investment decisions. For example, DoD Joint Chief of Staff Instruction CJCSI 3170.01D, March 2004, specifies a new Joint Capabilities Integration and Development System (JCIDS) that “implements a capability-based approach ... to identify improvements to existing capabilities and to develop new warfighting capabilities.”

To satisfy these guidelines, investment decisions cannot be made one system at a time, but must address how new candidate systems interact among themselves and complement currently fielded systems in order to achieve the maximum overall effectiveness of the resulting “system of systems” architecture. Thus the choice of an investment portfolio to support a system architecture must be made in the context of an overall mission framework. This approach is also consistent with the Clinger-Cohen Act (CCA)², which advocates a “portfolio investment process”. For example, the Air Force has a need to support a “Time-Sensitive-Targeting” mission by selecting the best combination of sensors, communication links, processing equipment, and weapon systems that, working together, can accomplish the mission most effectively.

What is PALMA?

The Portfolio AnaLysis MAchine (PALMA) is a tool that can be used to support the JCIDS process. Specifically, it is a computer program that helps decision-makers to select the best portfolio (combination) of investments from a set of potential investment options. The criteria for making the selection are usually cost (some other resource constraint could also be used) and mission-level performance. Some characteristics of PALMA are summarized below:

- Establishes an orderly process for approaching the investment decision.
- Is based on simple intuitive concepts.
- Forces identification of key objectives and capability gaps.
- Helps decision-makers to explain the reasons for their decisions.
- Promotes constructive discussion.
- Answers “what if” questions.

The key parts of performing an analysis with PALMA are (1) relating investment options to cost and detailed capabilities, and (2) relating detailed capabilities to higher-level goals. PALMA provides a framework for bringing this information together to identify

¹ PALMA is a trademark of The MITRE Corporation.

² The Clinger-Cohen Act (CCA), formerly known as the Information Technology Management Reform Act (ITMRA), Public Law 104-106, 10 February 1996.

the most promising portfolios of investments. In some cases the user may then wish to analyze the identified portfolios in more detail using other methods.

The Investment Portfolio Problem

Investment portfolio selection is analogous to the “knapsack problem”, as illustrated by the figure below. In words, given the limited space available in the knapsack (i.e., limited monetary budget), which items should a camper select to satisfy most of his needs for an upcoming hiking trip (i.e., to achieve maximum mission effectiveness)? Thus the camper must select those items that meet his most critical needs, maybe enough food and the first aid kit, and leave out those items that only satisfy “nice to have” desires, such as the camera. The process can get very complex as last minute items are “squeezed in”, or some items are removed to make room for a larger, but more critical item. Each possible combination of items that fit into the knapsack is a candidate “portfolio”. (Some combinations may not make sense because of the dependencies between items – e.g., don’t bring a flashlight without batteries, but PALMA can deal with this issue as well.)

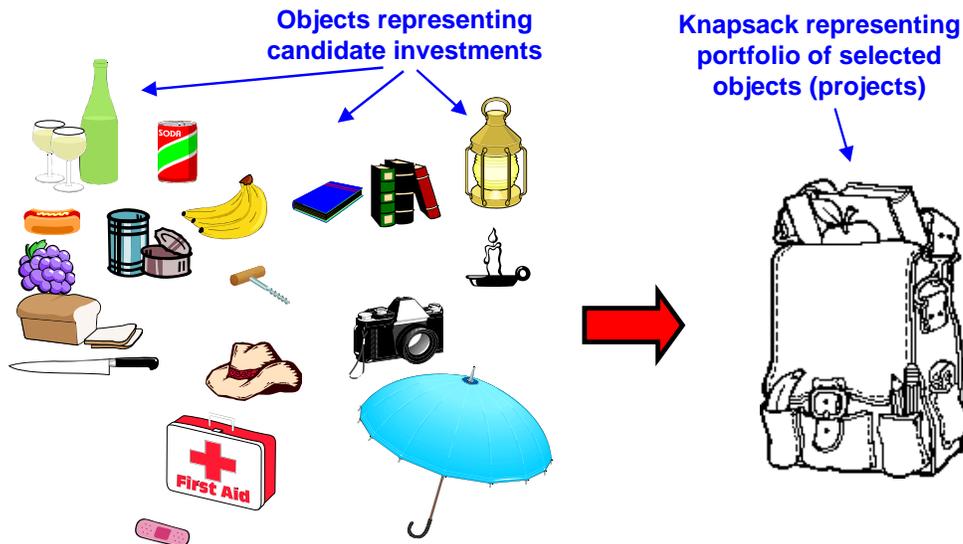


Figure 1. The Knapsack Problem

PALMA Features

The PALMA tool was specifically designed to help solve the investment portfolio problem described above. In particular, PALMA relates the impact of any set of investment options to the achievement of high level or “strategic” objectives. This tool can be used to support and automate the “strategy-to-task” breakout that is widely used in the military, or the “multi-objective” and “value-focused” models that are widely used in decision analysis. PALMA can be used in two ways: (1) to visualize the hierarchy of objectives and identify deficiencies in it, leading to the formulation of new investment options; and (2) to find the portfolios of investment options that are the most cost effective.

PALMA is founded on the principle that the process of making investment decisions should begin by finding answers to the following five questions:

1. What are our main mission objectives and what deficiencies do we currently have in capabilities that support these objectives?
2. What investment options should be considered to eliminate capability gaps?
3. How much do they cost?
4. What do they accomplish (in detailed terms)?
5. How do these detailed effects contribute to our overall goals?

PALMA provides a language in which the answers can be expressed and a mechanism for combining the information in a way that gives insight into the investment problem. Thus PALMA helps the user to draw conclusions from the information – especially which investment options should be funded to achieve the overall goals at the least cost.

PALMA Mission Tree Structure

To measure the impact of individual investment options, PALMA uses a strategy-to-task breakout that is constructed for the given mission (or overall goal) in the form of a tree hierarchy. In other words, the mission is split into its principal sub-missions, and then each of these sub-missions is further decomposed into subordinate missions, functions, and tasks. This approach, taken to its conclusion, results in a mission assessment model in the form of a tree with the top node representing the overall mission and the lowest-level nodes (leaves of the tree) representing the functions that constitute the basic elements of the mission.

The development of the PALMA strategy-to-task breakout is essential to measuring the mission impact of candidate investments. This endeavor does require a dedicated effort by mission area experts, but the PALMA tree can often be at least partially derived from other required studies, such as mission area analyses, concepts of operations, and architectural views. (The “operational views” or OV’s are particularly useful.)

Generally, a PALMA mission tree should be decomposed down to the level at which the contributions of individual investment options, i.e., the potential new fielded systems, can be directly evaluated. Thus, for a Time-Sensitive-Targeting mission, if one of the investment options is a new airborne sensor, then one of PALMA’s lowest leaf nodes should be “detect target”.

Figure 2 below gives a notional example of a PALMA strategy-to-task tree. In PALMA’s hierarchical or “tree” structure the most fundamental objective lies at the top of the hierarchy, shown in the figure as “targeting mission”. This overall objective is broken out into two lower-level objectives, “find target” and “attack target”, shown directly under it in the hierarchy. Each of these two second-level objectives is then broken down into subordinate functions at the next lower level. For example, in Figure 2 “attack target” depends on the functions of “plan attack” and “conduct attack”. Continued decomposition of a PALMA strategy-to-task tree can produce numerous nodes. PALMA provides an interactive graphics capability to facilitate the creation of a strategy-to-task tree. To help navigate through an extended hierarchy, PALMA also allows the user to view “sub-trees”, i.e., all the nodes in the PALMA tree that lie below a particular selected node.

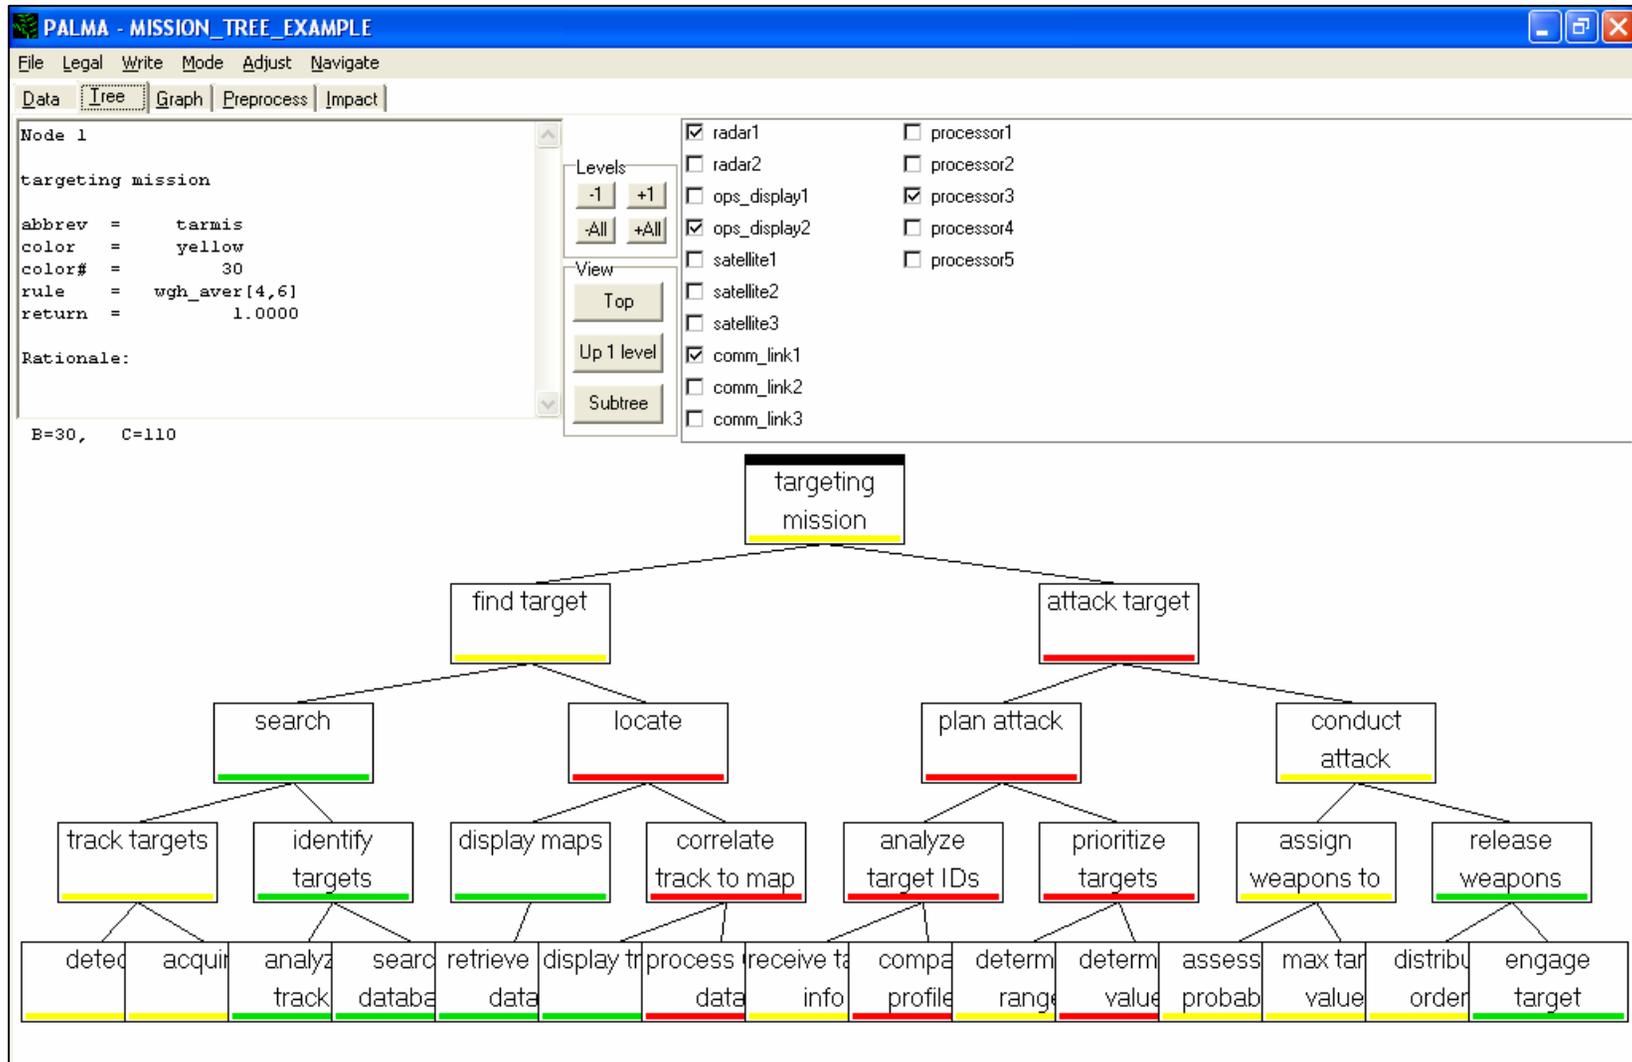


Figure 2. Sample PALMA “Strategy-to-Task” Tree

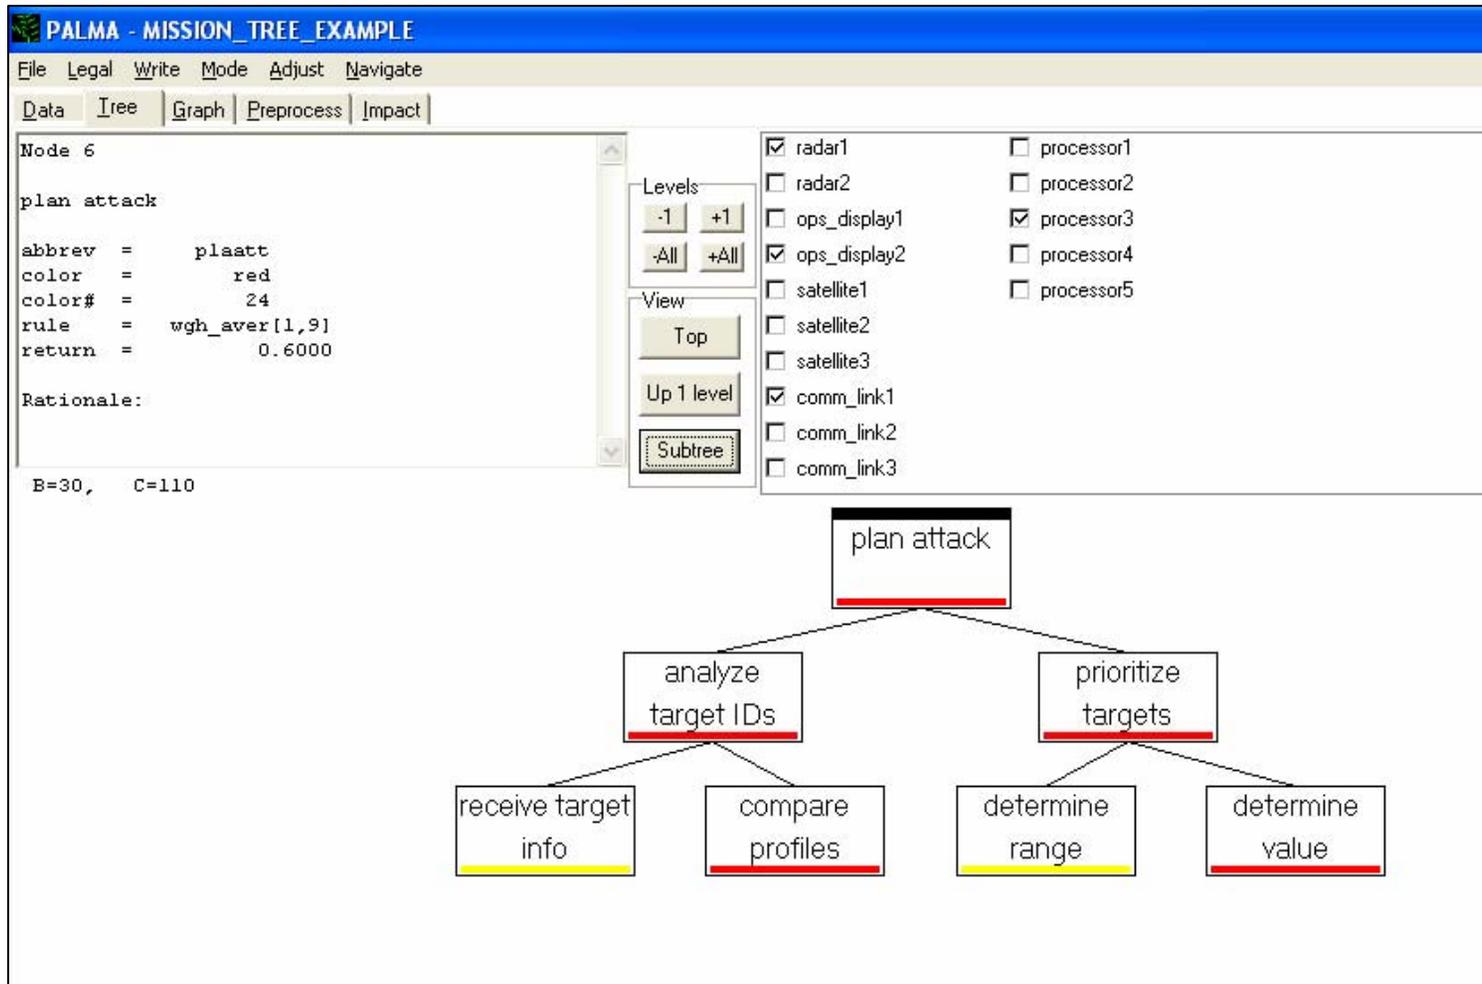


Figure 3. Sample PALMA Sub-Tree

For example, Figure 3 above illustrates the sub-tree headed by the “plan attack” function, which is broken out into “analyze target IDs” and “prioritize targets” tasks. The “analyze target IDs” task then consists of the more specific operations of “receive target info” and “compare profiles”, while the “prioritize targets” task then consists of “determine range” and “determine value” operations. These last four operations thus represent “leaf” nodes, where the impact of proposed investment options can most directly be determined.

Measuring Performance in a PALMA Tree

The performance of each node in a PALMA tree is expressed as a number in the range 0 to 100, representing the value to the expert mission area analysts (who helped develop the PALMA tree) and to top-level decision-makers. Performance scores are displayed in PALMA’s tree diagram by dividing the range 0-100 into smaller ranges and associating each range with a color. In our discussion we will use the familiar “stoplight” color system shown in Table 1. Thus, a node that is red in the tree diagram has a performance score from 0 to 25, which means that the node “Does not meet requirements.” Similarly, “Amber” (or “Yellow”) signifies a score from 26 to 50 meaning “Partially meets requirements,” “Green” stands for a score from 51 to 75 and means “Fully meets requirements,” and “Blue” represents a score above 75 and indicates “Exceeds requirements.” The color system may be changed for an individual analysis as long as it is applied consistently. PALMA uses the system below if none other is specified.

Table 1. PALMA Performance Color System

Color Name	Interpretation	Numerical Range
Blue	Exceeds requirements	75-100
Green	Fully meets requirements	50 - 75
Amber (or Yellow)	Partially meets requirements	25 - 50
Red	Does not meet requirements	0 - 25

The performance at any node in a PALMA tree depends on the performance score of its immediate lower-level nodes (its “children”). At each node above the leaf nodes, a “roll-up” rule determines how the node’s performance is computed from the performance of its children. The roll-up rule can be expressed by a variety of functions, e.g., a weighted average.

For example, in Figure 2 the display box at the upper left shows that the overall mission performance at the top “targeting mission” node is computed as the weighted average (wgh_aver[4,6]) of the performance of its two lower nodes, using a 40 percent weight for “find target” and a 60 percent weight for “attack target”. Also in Figure 3 the display box at the upper left shows that the performance of the “plan attack” function is computed as weighted average of its lower level task performance scores: “analyze target IDs” (10

percent) and “prioritize targets” (90 percent). PALMA provides a range of functions (min, max, etc.) that can be used for rollup rules.

Hence, once the performance scores of the leaf nodes have been determined, the performance scores of all higher level nodes can be computed in a “bubble-up” fashion using the specified rollup rules. This computation sequence determines the overall, top-level, mission performance numerical score and color.

The performance score specified for each leaf node is either the baseline score for that leaf or a better score induced by one of the selected investment options. The possible options are shown in the top right portion of the PALMA tree screen, as illustrated in Figure 2 above. Any of these options may be selected by the user by clicking in the associated box, and the impact on all affected leaf nodes will immediately be shown by a change in the colors of those nodes. The rollup rules will also then reflect the impact on higher nodes via a “ripple effect” up the tree. In particular, Figures 2 and 3 show the effect of selecting the four investment options; “radar1,” “ops_display2,” “comm_link1,” and “processor3.” In addition, the cumulative total cost (C) and benefit (B) (the performance at the top node), resulting from all investment options chosen so far, will be shown just below the upper left display box. (e.g., B=30 and C=110 in Figure 2). (Note: All costs and benefits shown in our examples are purely notional.)

By examining the colors in the tree diagram, the user can quickly identify the relative strengths and deficiencies among the various areas of mission performance. This feature can be very helpful in determining the need to formulate new investment options to meet unsatisfied requirements.

For example, in Figure 2, the overall “targeting mission” is shown as “Yellow.” Since this node is highlighted (indicated by the black bar across the top), the upper left display box shows that its numerical score is 30 and its score is calculated via a weighted average “wgh_aver[4,6]” roll-up rule. Thus, the “targeting mission” performance is calculated as a weighted average of the scores of its two children, with the second child, “attack targets,” specified as having a weight of 6 out of the total of 10 (4+6) – and so receives 60 percent of the weight. Clearly then, the most efficient way to improve the “targeting mission” performance is to increase the performance score of the “attack targets” objective. Tracing further down the tree, we next see that one child that “attack targets” depends on is “plan attacks,” which is currently rated as “Red” and whose sub-tree is shown in Figure 3. From the roll-up rule, “wgh_aver[1,9],” shown in the upper left display of Figure 3 we can see that “prioritize targets” is the most important task below “plan attacks.” Finally, we can trace down to find two (equally weighted) leaf nodes below the “prioritize targets.” Since one of these, “determine value,” is currently evaluated as “Red,” we may wish to improve its performance by selecting additional investment options. The resulting impact will then “bubble up” the PALMA tree to improve the score of our overall SAR Mission performance.

Figure 4 again shows the “targeting mission” PALMA tree, but with an additional investment option, “processor5,” selected (as indicated by the check in its box). Also note in Figure 4 that the “determine value” leaf node has been highlighted, so that the

upper left display now shows which investment options affect that task, and what performance score each would induce. Note that “processor5” is the only investment listed and turns the “determine value” node to “Green” with a numerical score of 63. By comparing Figure 4 to Figure 2, also note that six other leaf nodes have had their performance improved by this selection – some have been turned to “Green” and others to “Blue.” Finally, note that Figure 4 now reads “B=52, C=280” just below the display box, so that the top level “targeting mission” performance has been improved to (low) “Green” – at an additional cost of 170 (hypothetical) units. Thus the investment “processor5” has a high impact, but also bears a high cost.

Thus, the PALMA user can successively select investment options in this “trial and error” fashion to try to improve the performance of as much of the strategy-to-task tree as possible within an allowable budget. However, PALMA provides a much easier way to find the most cost-effective investment portfolios.

PALMA Optimization of Investment Portfolios

The PALMA tool automates the process of determining the most cost-effective portfolio of investment options at any given budget level. When there are a modest number (25 or fewer) of investment options, PALMA performs an exhaustive search of all possible investment portfolios. PALMA evaluates the respective total costs and top-level mission benefits of these possible portfolios and displays the results on the PALMA Graph page as shown in Figure 5 – where Benefit is shown on the vertical axis and cost on the horizontal axis.

The points on the upper left edge of this graph are referred to as the “efficient frontier” of possible portfolios. These points are highlighted to reflect that fact that the portfolios they represent produce the highest possible benefit for the associated cost. The actual portfolio that any such point represents is immediately displayed (by the checked boxes) when that point is selected, together with its associated cost and top-level mission benefit. For example, in Figure 5 the selected point (shown in red on the efficient frontier) represents the portfolio; satellite1, comm_link2, and processor5, which has a total cost of C=220 and mission benefit (performance score) of B=63 (mid-Green). Thus, for a cost of 220, there is no other portfolio of investments that produces a higher benefit than 63.

Also note that there are three labels, “Un”, “In”, and “Out”, next to each investment option in Figure 5. These labels allow the user to place constraints on the choices that PALMA explores in constructing the efficient frontier. Specifically, “Un” means that the choice of that investment option is unrestricted (i.e., it may be chosen or not), “In” means that the investment option must be included within any portfolio chosen, and “Out” means that the investment option must be excluded from any portfolio chosen. This functionality allows the PALMA user to address extraneous issues (e.g., political considerations) and perform “what-if” analyses.

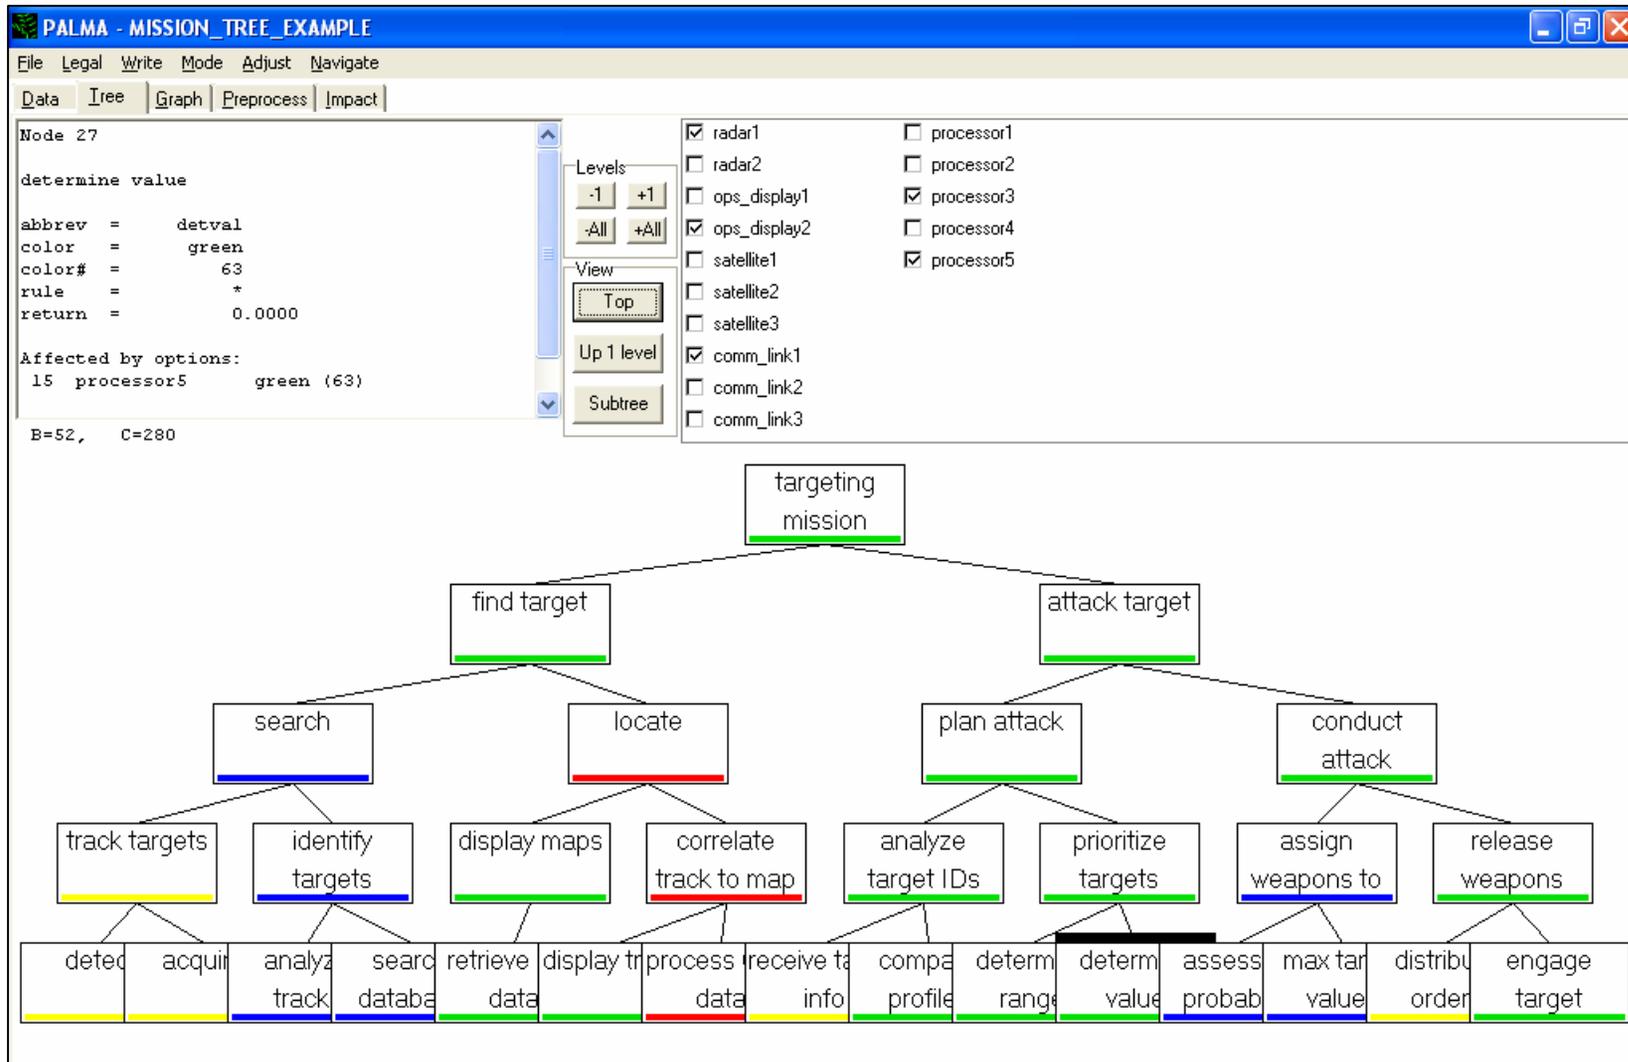


Figure 4. PALMA Tree with Added Investment Option

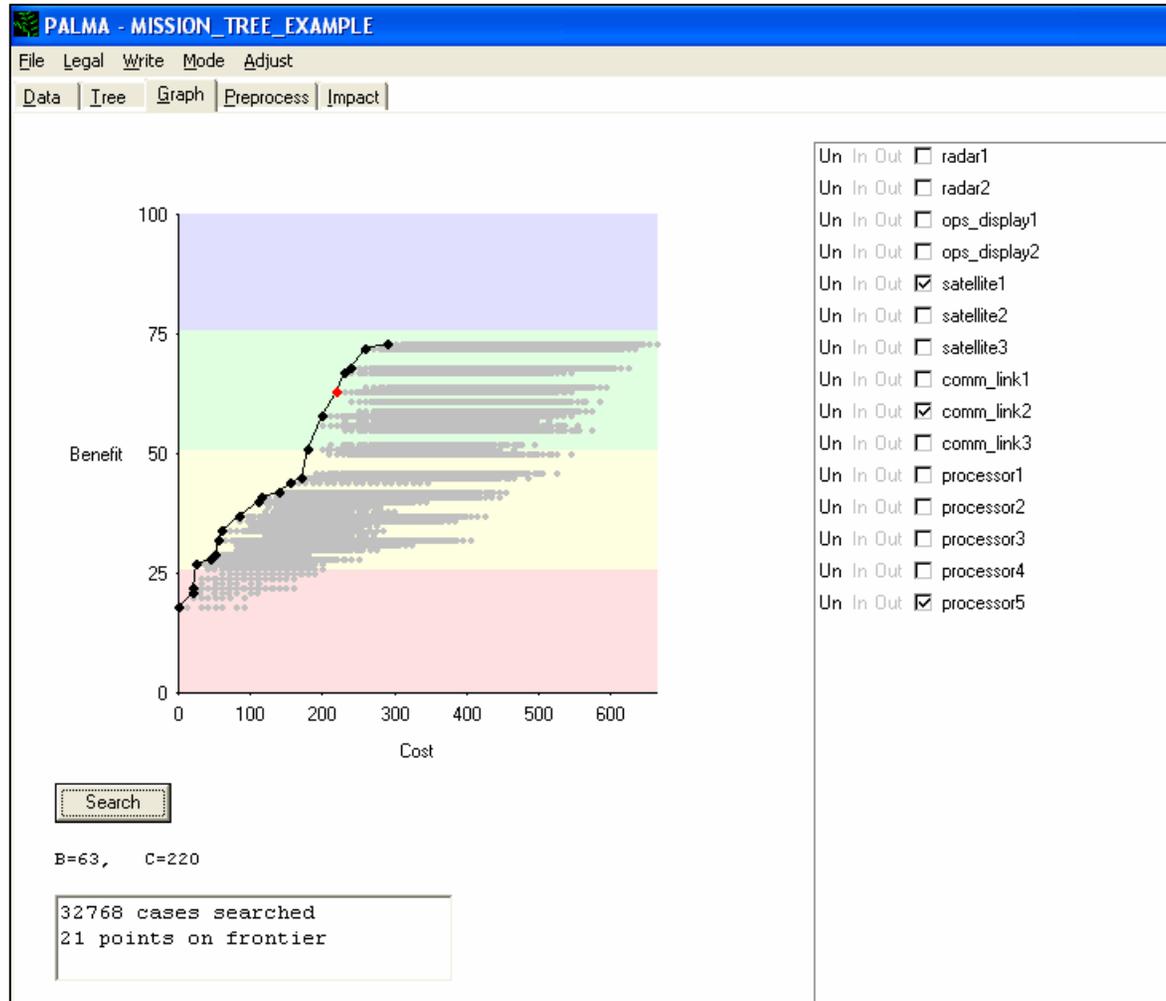


Figure 5. PALMA “Efficient Frontier” Graph Page

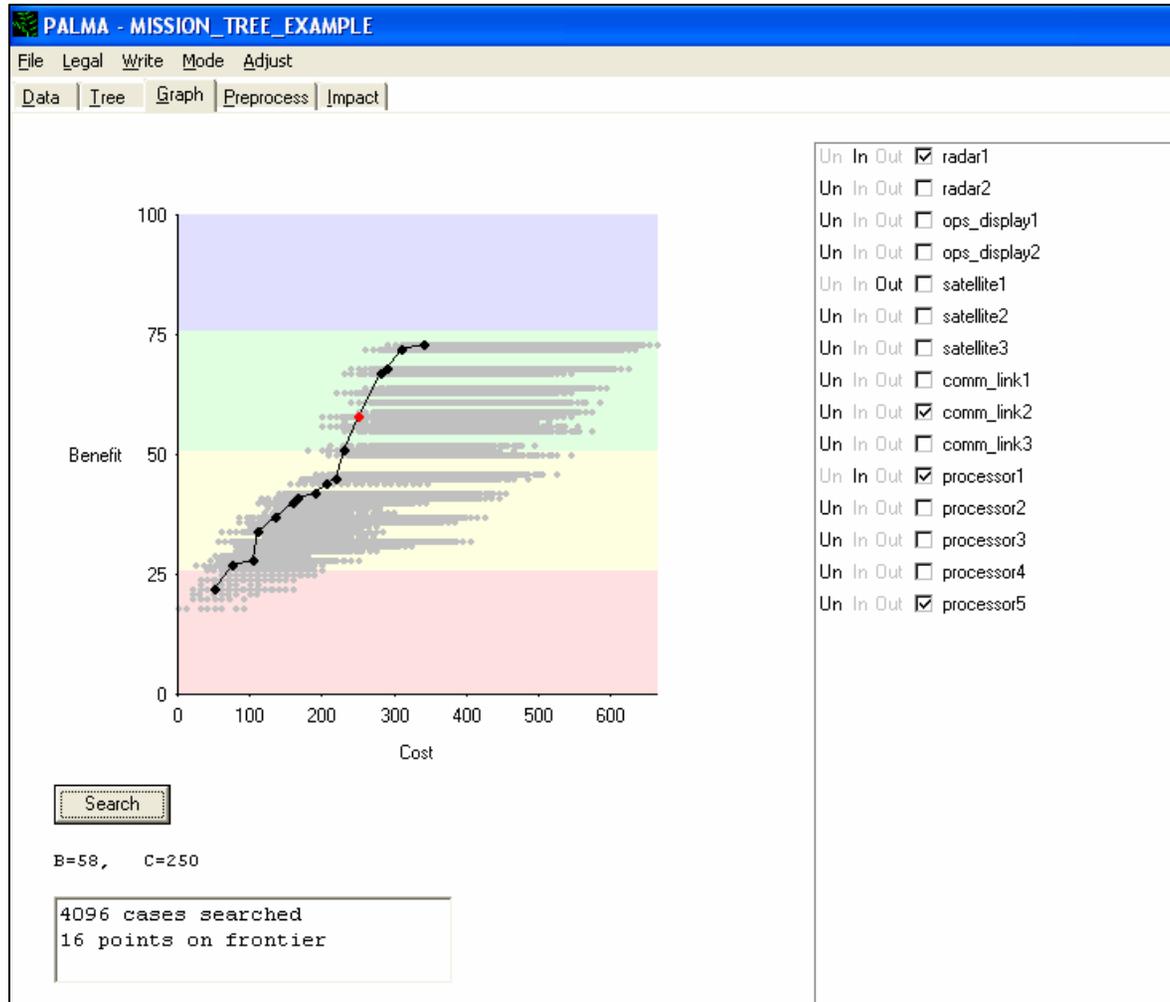


Figure 6. PALMA “Efficient Frontier” with Restrictions

For example, Figure 6 above illustrates how the efficient frontier is shifted to the right as a result of imposing the restrictions of forcing *in* the options “radar1” and “processor1”, and forcing *out* “satellite1”. Thus the selected point on the new efficient frontier shows a benefit of 58 – lower than the benefit provided by the selected point on the original efficient frontier shown in Figure 5. However, these restrictions have increased the cost of achieving this benefit level to 250 – higher than the previous 220 cost required to achieve the higher benefit of 63 without the restrictions.

The PALMA tool can actually address 100 or more investment options by modifying its search method from an exhaustive enumeration to operations research optimization techniques, such as a “genetic” algorithm or integer programming. In addition, PALMA offers several other useful features. The most notable of these capabilities include the ability to address dependencies among options, and the ability to plan and manage evolving portfolios over a period of time increments while subject to multiple budget constraints (e.g., by year and by budget category).

To address dependencies among options, PALMA can flexibly characterize relationships (dependencies) among potential investments. Dependencies provide a way for the user to restrict which combinations of options are allowed in a portfolio. These dependencies include both ‘required combinations’ (where selection of one option implies the necessity of selecting an additional option) and ‘excluded combinations’ where selection of an option excludes selection of other specific options (e.g., versions). In addition, the user can apply PALMA’s substantial selection of ‘roll-up rules’ to address more complex relationships among investment options.

To support PALMA’s utility to enhance the planning and management of evolving portfolios over time while subject to multi-year budget constraints, PALMA also is available in a time-phased version. This multi-year version³ considers not only *which* options to select for the portfolio, but also *when* to select them. PALMA searches for the most cost effective portfolios based on the multi-period benefits and multi-period costs. Two interpretations of multi-period cost are supported:

- a. The total cost over all periods and cost types, or
- b. The maximum over all periods and cost types, in terms of the ratio of cost to budget for that (period, cost-type) pair.

These approaches allow us to find portfolios that have low overall cost or that come as close as possible to staying within the budget in each period and cost type. This is supported by the evaluation of costs and benefits where costs for each option are specified by a time phase (discrete time phases must be identified), and type of cost (‘color of money’) subject to the constraints of a budget which is also specified by time phase and type of funding. Option impacts are also specified by time phase.

³ Multi-year version of PALMA is described in greater detail in “Portfolio Analysis Machine User's Guide – Multi-Year Version 2.19,” by Brian K. Schmidt, April 2005 (Draft MTR).

Summary

PALMA provides a user-friendly, PC-based, interactive environment for identifying optimal investment portfolios in terms of the overall “system-of-systems” mission capability that they provide. Its strategy-to-task visual display helps the user to identify the strengths and weaknesses of his current capabilities in terms of essential mission functions and tasks. PALMA then provides the user with two different approaches to reduce shortfalls in capability. He can either manually select those investment options that can eliminate specific deficiencies, or he can examine the efficient frontier graph that automatically specifies the most cost-effective portfolios across the range of funding levels.

Of course, a PALMA analysis should be tied in with other planning methods (models, simulations, analysis of requirements, architectures), and these sources can provide many of PALMA’s required inputs. In addition, by exposing remaining “gaps” in mission performance, a PALMA study may suggest new investment needs or combinations of investments to be studied further.

Thus, the PALMA analysis approach provides the investment decision-maker with a systematic means for establishing a substantive and defensible “capability-based” rationale for investment recommendations.

Contacts:

MITRE Bedford, MA:

Brian Schmidt, bschmidt@mitre.org (781-271-6186)

Dick Moynihan, moynihan@mitre.org (781-271-6227)

Elaine Goyette, esg@mitre.org (781-271-6031)

Bruce Lamar, bwlamar@mitre.org (781-271-8208)

MITRE McLean, VA:

Tony Shimi, tshimi@mitre.org (703-883-7391)