Technical Report

Department of Computer Science
and Engineering
University of Minnesota
4-192 EECS Building
200 Union Street SE
Minneapolis, MN 55455-0159 USA

TR 01-012

Extending the Product Family Approach to Support n-Dimensional
and Hierarchical Product Lines

Jeffrey M. Thompson and Mats P. Heimdahl

February 26, 2001

Extending the Product Family Approach to Support n-Dimensional and Hierarchical Product Lines*

Jeffrey M. Thompson and Mats P.E. Heimdahl
Department of Computer Science and Engineering
University of Minnesota
4-192 EE/CS; 200 Union Street S.E.
Minneapolis, MN 55455 USA
+1 (612) 625-1381
{thompson,heimdahl}@cs.umn.edu

Abstract

The software product-line approach (for software product families) is one of the success stories of software reuse. When applied, it can result in cost savings and increases in productivity. In addition, in safety-critical systems the approach has the potential for reuse of analysis and testing results, which can lead to a safer systems. Nevertheless, there are times when it seems like a product family approach should work when, in fact, there are difficulties in properly defining the boundaries of the product family.

In this paper, we draw on our experiences in applying the software product-line approach to a family of mobile robots as well as case studies done by others to (1) illustrate how domain structure can currently limit applicability of product-line approaches to certain domains and (2) demonstrate our initial progress towards a solution using a set-theoretic approach to reason about domains of what we call n-dimensional and hierarchical product families.

1 Introduction

Today's consumers want devices and systems custom fitted to their needs. Software product-line engineering has the potential to deliver great cost savings and productivity gains to organizations that provide families of products, as well as give those organizations a competitive edge in the market-place. For safety-critical systems, software-product line engineering has the potential to produce systems that are more safe than their serially produced counterparts while being cheaper and faster overall to build.

Although one of the main barriers to the use of product family techniques is one of process and orga-

nizational acceptance, technical issues have not been completely solved for product-line engineering. The techniques available work best for cohesive product families, where the variabilities do not have complex interdependencies. When this is not the case, it can be difficult to apply the product family approach even though there might be significant commonalities between the members of the family. As an alternative, we propose to view the families themselves in a multi-dimensional and hierarchical fashion. This helps us to deal with existing problems, for example, *near commonalities*, and also, helps to extend the approach to domain which, traditionally, would be difficult for product-line engineering.

The paper is organized as follows. The next section presents background on work in product-line engineering. Section 3 presents current issues with the product family approach. Then, we discuss the foundations of our approach in Section 4. Section 5 presents how our approach can be used to address existing issues in product-line engineering. The mobile robotics case example is in Section 6. Finally, Section 7 presents our conclusions and ideas about future directions for this work.

2 Product-line engineering background

The notion of a product family was introduced by David Parnas in [16]. According to Parnas, it is desirable to study a group of programs as a whole whenever the programs share extensive commonalities. Parnas observed that often programmers would create new programs by modifying existing programs. This process usually involved a reverse step where parts of the working program were discarded. Furthermore, the new program was sometimes crippled by design assumptions made for the original program which did not apply to the new program. Thus, Parnas postulated that it would better to start out by defining what was common about all such programs and successively refining the design until you had working programs as

*This work has been partially supported by NSF grants CCR-9624324 and CCR-9615088, and by NASA grants NAG-1-2242 and NCC-01-001.

Figure 1. A simple product family

the leaves of a tree structure, with nodes within the tree representing the various design decisions.

Batory and O'Mally [3] discussed how to reuse large portions of a system based on breaking it into components and introduced a simple language for describing the components and their composition. Gomma [10] discusses using domain modeling [17] to create a centralized library of components which are then used by a generation facility to produce the target application.

Weiss and Ardis [19, 2] developed the FAST (Family-oriented Abstraction, Specification and Translation) process that integrates the above with specialized languages [15, 4] for each domain. A similar process is mentioned by Campbell *et al.* in [8, 7] and also by Lam [13, 12]. The differences between these works are primarily in the sort of artifacts produced by domain engineering.

The commonality analysis [18] is a central feature of product-line engineering. This is the document that notes all the commonalities, i.e., features which are present in all systems in the domain, and variabilities, i.e., features which distinguish the different members of the domain. The commonality analysis defines the requirements for the product line.

One way to view a product family is as a set, where the boundaries of the set are determined by the commonalities, and the individual members of the set are distinguished by the values of their variabilities (Figure 1). As the figure demonstrates, it is entirely possible that some members of the family may theoretically exist but not yet be built (shown in gray). Furthermore, the family may be undefined at some points within the boundaries due to, for example, illegal or nonsensical combinations of variability values. We will use this view of a product family throughout the paper to demonstrate how current approaches to product-line engineering might be expanded to a greater class of systems.

3 Extending product families

Current techniques for product-line engineering work well if the following conditions are met:

- The systems in the family share significant commonalities, and
- The variabilities which define each family member have a straightforward decision model, i.e., it does not require many complicated rules to de-

scribe how the variability values are assigned to produce each family member.

The first point describes the essential feature of product families that Parnas noticed in his work. However, the second point originates in the practical experience of many researchers who have labored to construct software product-lines. Robyn Lutz observed that the primary limitations of the product family approach stem from difficulties in handling “*near-commonalities* and relationships among the variabilities”[emphasis added] [14]. Thus, the more simple the relationships among the variabilities, the easier it is to construct the product family.

3.1 *n*-Dimensional product families

Attempts have been made to organize the product family requirements in a hierarchical fashion [14, 16, 11, 12]. Lutz noted in her attempt to organize the variabilities into a tree that “there were several possible trees, with often no compelling reason to select one possible tree over another” [14].

Brownsword and Clements present a shipboard command and control systems family which contained 3000-5000 parameters of variation for each ship [6]. They state that “the multitude of configuration parameters raises an issue which may well warrant serious attention.” In addition, they present three different views of the architectural layering of the base system that “do not conflict with each other; rather they provide complementary explanations of the same ideas.”

Both these examples, as well as our own experience in the domain of mobile robotics, illustrate the fact that often a product family is multi-dimensional; therefore, a hierarchical decomposition is not sufficient to capture the structure of the domain. We call such domains *n-dimensional product families*.

3.2 Hierarchical product families

Suppose that a company wished to construct a flight guidance system (FGS) for both fixed-wing aircraft and helicopters¹. The FGS is responsible for issuing commands that keep the aircraft level, cause it to climb or descend, and so forth. Furthermore, the FGS must interact with other airborne systems. Many of the tasks that the system has to perform might be common across these two radically different aircraft: interaction with other systems, deciding to level off at a particular altitude, mode transition logic related to when it is legal to switch between the various operating modes. Therefore, many requirements between these two systems will be the same, or very similar. Nevertheless, the actual control of the aircraft is very different. Therefore, developing a single set of commonalities and variabilities which span this entire domain is difficult.

¹We would like to thank Steven P. Miller of Rockwell-Collins Inc. for this example

Figure 2. FGS product family covering flying craft

Some would argue that this difficulty stems from the fact that the family is simply too diverse to be considered a product line. However, it is clear that these systems share much in common, which was the original, and in our view the most important, criterion for being a family. Thus, we propose the concept of a *hierarchical product family*.

Most previous attempts at product family structuring have focused on hierarchically grouping the *variabilities* while the *commonalities* remain the same for all family members [14, 12]. Notable exceptions are Parnas [16] and Brownsword and Clements who noted in their case study at CelciusTech [6] that sometimes product-lines exist within the main product line. However, Parnas' work is based heavily on design and coding choices; therefore, it is difficult to apply to requirements. Brownsword and Clements mention this phenomenon in passing and apply it in a more limited way than what we advocate.

In our approach, *additional commonalities* which are *unrelated* to the parent product family can be added in the sub-families. The hierarchical decomposition of the FGS family is shown in Figure 2. Thus, the helicopter sub-family can have significantly different requirements than for fixed-wing aircraft, yet share many things in common as well.

This will eventually effect the architecture and structure of the systems. For example, the product of the domain engineering for the parent family, Flying Craft FGS, might be a set of reusable components, whereas the product of domain engineering for the children might be a reference architecture or generation facility. The architectures for the fixed-wing aircraft and the helicopters could differ significantly and use the components from the parent family in different ways.

By structuring the requirements in this way, we have avoided imposing restrictive design constraints on the family members and instead focus on the structure of the domain itself. Furthermore, should the company wish to start building FGS systems for an entirely new set of aircraft, for example, tilt-rotor aircraft, this could be done while reusing many aspects of the FGS systems already implemented. This is also shown in Figure 2.

3.3 Constraints on the solution

When starting to develop a structuring technique for product family requirements that would be able

Figure 3. Hierarchical decomposition and subset structure

to deal with n-dimensional and hierarchical product families, we determined that any such structuring technique must:

- Be based on structures that are present in the *domain* itself, not on implementation or design concerns,
- Be simple, allowing the analysts to capture the structure of the domain without introducing complex notations or concepts,
- Be amenable to the types of structures observed in product family analysis, and
- Produce a readable and usable artifact that facilitates reasoning about the structure of the domain.

We chose to explore a structuring technique based on a set-theoretic view of product families. This notation proves surprisingly useful for thinking about the structure of a software product line, yet is simple and based upon well understood principles.

4 Structuring technique

As mentioned previously, a product family can be described in terms of a set, where the boundaries of the set are defined by the commonalities and the members of the set are distinguished by the values of the variabilities. This section describes how set theory can be used to think about structuring product families.

The most basic structure that can be represented with the set theoretic approach is the subset. Figure 3 shows a product family, **A**, which has been divided into two subsets, **B** and **C**. Furthermore, **C** has been further divided into subsets **D** and **E**. This corresponds to a hierarchical decomposition of the family.

Consider a member of family **E**, e_1 . The member e_1 must have all the commonalities defined for **E** as well as have some values for all the variabilities in **E**. Furthermore, because **E** is a subset of **C** and **A**, e_1 is also a member of families **C** and **A**. The general definition for any family **E** which is a subset of another family **C** is as follows:

- **E** must include all of the commonalities in **C**.

Figure 4. Abstract versus non-abstract families

- **E** must include all of the variabilities in **C**; however, **E** may restrict the range or options available in the variabilities.
- **E** can add additional commonalities which are not present in **C** as long as the additional commonalities do not conflict with the commonalities or variabilities in **C**. These new commonalities might come from a refinement of variabilities in **C** or might be completely unrelated.
- **E** can define additional variabilities which are not present in **C** as long as those variabilities do not conflict with the above.

The first criterion is straightforward and necessary for the subset **E** to be completely contained within **C**. The second criterion defines the fact that **E** may wish to refine or restrict the values of the variabilities of **C**. For example, in the mobile robotics domain, a variability across the entire domain might be that the maximum speed of the mobile robot can vary from one to five miles per hour. However, subsets might define a lesser maximum speed depending on the hardware involved. It is possible for this refinement to result in an additional commonality, for example, a subset that instantiates a boolean choice variability to a particular value. Additional commonalities can also be added which are unrelated to the parent family. For example, it is likely that the family of helicopters will need different commonalities than the family of fixed-wing aircraft. Finally, it is possible to add additional variabilities.

The two cases of hierarchical decomposition are shown in Figure 4. Part (a) of the figure demonstrates that the family **R** need not have any members that only exist in **R**. In a sense, **R** is an *abstract family*, because any member of **R** must be either a member of **S** or a member of **T**. This is similar to our FGS example from earlier, where all family members are either helicopters or fixed-wing aircraft and it does not make sense to talk about member which are only of the parent family. However, this need not be the case, as Figure 4(b) demonstrates. In the mobile robotics domain (see Section 6), we will have a basic robotic platform which will form the outer family member. This outer family will *not* be abstract because there are some robots which only conform to the minimum specification.

Another structure that can be represented using a set-theoretic approach is that of set intersection. The ability to represent a set intersection distinguishes this

Figure 5. Set intersection and non-hierarchical structure

approach from the purely hierarchical structures which have been applied by others. This is shown in Figure 5.

Consider a member, m_1 of **M**. By definition, m_1 is also a member of families **K**, **L**, and **J**. Thus, m_1 must have all the commonalities of both **K** and **L**. In addition, **M** is a subfamily of both families **K** and **L** (this is shown in the figure). The constraints on any family **M** which is a subset of families **K** and **L** are as follows:

- **M** must include all the commonalities of both **K** and **L**.
- **M** must include all the variabilities of both **K** and **L**; however, it may restrict those variabilities as above for subsets.
- **M** may introduce additional commonalities which are not present in either **K** or **L**.
- **M** may introduce additional variabilities which are not present in either **K** or **L**.

These structures can be used to document and reason about the two problems explored in the previous section; they can be used to describe product families which are both *n-dimensional* and *hierarchical*.

5 Addressing existing issues

This section describes how our approach can assist with well-known documented issues in product-line engineering. We describe how our structuring method can deal with near-commonalities as well as variability dependencies.

Near-commonalities: A near-commonality (NC) is a commonality which is true for almost all (e.g., all except one) member of the product family. Lutz states that in her experience near commonalities “frequently had to be modeled” [14]. One solution for near commonalities is to model them as variabilities; however, this is, in some sense, a misrepresentation of their basic properties. The solution that Lutz advises is to model it as a constrained commonality of the form “If

Figure 6. Set representation of a near-commonality

not member n then NC_1 .” However, a complex domain might contain numerous constrained commonalities with conditions significantly more complex than the example just mentioned.

Figure 6 shows how a near-commonality is represented in our approach. The near commonality, NC_1 , would simply be a property of family Q (and not of P). Thus, the commonality naturally does not apply to n , a member of only P but does apply to any member of Q . This has several advantages. First, NC_1 is now a pure commonality of Q . Second, if another member of the family is introduced with reduced functionality [14] it need only be added as a member of P and Q may remain untouched. Finally, the subset structure can act as a guide in determining that certain components in the eventual application engineering environment will not be needed for n .

Dependencies among options: In [14], Lutz cites modeling dependencies among options as one issues that must be addressed in product family engineering effort. A dependency is typically a constraint among the variabilities, for example, if variability V_1 has value B then variability V_2 must have option C . Ardis recommends treating this constraint as a commonality. However, in our experience, without some additional structuring, the domain could become littered with such commonalities; in addition, it may not be clear given a set of constraints whether or not a particular variability is viable.

In our approach, we can also represent constraints like these as commonalities. However, we isolate them into logical groups by forming different subfamilies so that their numbers do not become overwhelming. In the abstract example given above, a set would be defined where “ V_1 has option B ” and “ V_2 has option C ” are both commonalities.

In this section, we have discussed how our approach can help deal with existing issues which have been raised in the literature regarding product families. In the next section, we go on to describe how this structuring mechanism can help deal with more difficult product families illustrated with an example in the mobile robotics domain.

6 Mobile robotics: An n -dimensional and hierarchical example

The domain of mobile robotics that we have examined encompasses small robots ranging in size from approximately 6 inches long to up to about 2 feet long. Also, the robots can support many different behaviors – scouting an area, constructing a map, working collectively, and so forth. We wished to model the mobile robotics domain as a product family taking into consideration both the different hardware platforms that could be supported and the many different behaviors we wished to specify. This proved to be difficult using conventional product-family techniques because the mobile robot family is both n -dimensional and hierarchical [9].

The mobile robotics domain breaks down along two clear dimensions: the hardware platform and the desired behavior. Each hardware platform conforms to a basic specification: it can move forward and backward, turn left and right, sense whether or not an object is in front of it, and so forth. In addition, the hardware platform may or may not be equipped with some sort of vision system or infra-red camera; the various sensors used to monitor the environment differ greatly in the speed and accuracy with which they provide information. On the behavior side, we can imagine that a basic behavior might be a random exploration of the robot’s environment where the primary goal of the robot is collision avoidance and recovery. Furthermore, more complex behaviors can be added, for example, wall following, going through doors, and finding particular objects. Therefore, along both the hardware and behavior *dimensions*, the mobile robot family can be viewed *hierarchically*.

Of course, there are constraints between the two dimensions and not all behaviors can run on all hardware platforms. For example, a behavior that requires the robot to find all red objects in a room will not work unless the robot has a sensor capable of distinguishing red objects from non-red objects.

As a specific example for this paper, we will consider a mobile robotics domain consisting of three classes each of robots and behaviors. Although this is a *significant* simplification of the actual domain, it will be sufficient to illustrate the benefit of the structuring techniques that we are proposing.

A family member will consist of a pairing of the desired behavior with the robotic platform. We will first explain each dimension in some detail and then look at an overview of the family composed of these two dimensions. As a notional convention in the following two sections, commonalities and variabilities which deal with the hardware are noted [C_H] and [V_H] respectively and those which deal with the behavior are noted [C_B] and [V_B].

6.1 Hardware dimension

Along the hardware dimension, we will consider a limited subset of the robot domain containing three families of hardware. The actual domain is much more

complex; it includes many more types of sensors and different actuators, for example, a gripper or robotic arm that can be used to pickup and move objects in the environment. We will consider the following three classes of mobile robotic hardware in this paper.

1. A basic robot with forward and backward motion capabilities, a range sensor that give distance to the nearest obstacle and whether or not the obstacle is on the right or on the left, and a forward collision detection mechanism.
2. The basic robot with the ability to distinguish between obstacles that are straight ahead versus only the right or left (i.e., better granularity in the estimation of the obstacle's position).
3. The basic robot with the ability to distinguish the color of objects in its environment.

The following paragraphs describe the properties of the various hardware platforms that we will consider.

Basic platform: A basic feature of our robotic platform will be that it can move around its environment in some fashion. Thus a common feature of the robots is the following:

C_H1.1 Each platform will provide a basic means of locomotion; it will have the ability to turn a specified number of degrees from the initial heading, move forward, move backward, and stop.

Nevertheless, the robotic platforms that we will consider differ greatly. Some platforms are commercially built whereas others are built in-house, for example, out of Lego building blocks and small motors. The following variabilities capture these ideas:

V_H1.3 The hardware comprising the robotic platform varies

V_H1.3a The means of locomotion may vary (e.g. treads, wheels, legs, etc.)

V_H1.3b The maximum speed of the robot varies.

V_H1.3c The control of locomotion varies. The locomotion system may provide simple on/off values or real or digital valued representation of speed and direction.

V_H1.3d The type of input expected by the locomotion system varies. It may expect boolean, real, or digital values indicating speed and direction of the platform.

V_H1.3e The size of the platform varies. This will dictate the amount of room needed to turn or avoid an obstacle.

In order to avoid running into obstacles in the environment, the robot must have some kind of range finder. The platform must also be able to tell whether or not the obstacle is on the right or left so that it can

take actions to avoid hitting the obstacle. However, range finders vary significantly in the type and quality of information they provide. For example, a sonar sensor provides a wide field of detection but is noisy and inaccurate. A laser range finder, on the other hand, will provide distance with high accuracy and can detect even small obstacles.

C_H1.2 All platforms will have at least one range finder that will provide input to the system regarding the detection of an obstacle.

C_H1.2a The range finder will provide an indication of the distance to the obstacle.

C_H1.2b The range finder will provide an indication of the location (right or left) of the obstacle in relation to the robot.

V_H1.1 The number and type of devices used for range finding is likely to vary. The type of output generated by the range finder varies. Different range finders may provide output as a real-valued estimate, a digital estimate, or a boolean indication of obstacle detection.

Finally, because the mobile robots operate with such noisy and inaccurate sensors it is a certainty that they will occasionally have collisions. Thus, platforms must have a method of detecting collisions so that they can perform recovery actions in the behaviors. This could be implemented in a variety of different ways, for example, by installing bumpers on the robot or by detecting that the motors that drive the wheels have stalled.

C_H1.3 All platforms will have at least one mechanism for detecting collisions.

V_H1.2 The number and type of collision sensors(s) varies and the type of output generated by the collision sensor varies.

Enhanced obstacle detection: Some platforms may have more advanced sensors to detect obstacles. For example, a robot with an array of sonar sensors arranged in an arc can get much more information about potential obstacles than merely whether they are on the right or on the left. For enhanced obstacle detection, the robotic platform should be able to detect whether or not it has an obstacle in front of it in addition to obstacles on the right and left.

C_H2.1 Platforms will have the ability to distinguish whether an obstacle exists directly in front of them as well as whether it is on the right or on the left. See related [C_H1.2b]

V_H2.1 The granularity of obstacle position detected will vary. For example, some platforms may provide an enumerated indication of left, right, or front for the obstacle whereas some may provide an estimated degrees to the obstacle.

This sensing capability allows the robot to perform more complex behaviors, for example, maneuvering closer to obstacles or going through doors.

Environmental vision: Some robots may be equipped with a camera or other sensing device that can give them information about the color objects in their environment. The type and quality of robotic vision systems varies greatly; however, most can distinguish between primary colors.

C_H3.1 Platforms will have a sensor capable of determining the color of objects in their environment; for example, the sensor should be able to distinguish between red objects and blue objects.

6.2 Behavior dimension

The behavioral dimension defines *what* the robot does. Of course, the behavior of the robot is highly related to the hardware dimension, which constrains what the robot *can* do and what information about the environment is available. Nevertheless, to a large extent the behaviors can and should be reused across different hardware platforms. The spectrum of behaviors possible, even with the limited hardware classes that we have defined, is large. For the purposes of this report, we only have space to discuss a few of them. Thus, along the behavioral dimension, we will consider the following classes of behavior.

1. Random exploration, where the robot moves around its environment attempting to avoid obstacles.
2. Random exploration with the ability to negotiate doors.
3. Random exploration with the ability to signal when it encounters objects of a particular color.

The following paragraphs discuss these behaviors in more detail.

Random exploration: Rodney Brooks [5] recommends a layered architecture of robotic behaviors. Our approach to modeling the behavioral dimension is similar in that our basic behavior is a random environmental exploration. While exploring, the robot should attempt to avoid obstacles in the environment.

C_B1.1 The robot shall attempt to avoid colliding with obstacles in its environment using its sensors to detect obstacle(s) and changing its course or speed to avoid the obstacle.

V_B1.1 Although detected by the robot's sensors, an object may or may not be considered an obstacle depending on the robot's mode of operation. See, for example, [C_B2.1a]

As mentioned previously, because of the robot's noisy and inaccurate sensors it is likely that the robot will sometimes collide with an obstacle. When this occurs, the robot should attempt to recover from the collision and continue exploration.

C_B1.2 If the robot collides with an obstacle, it shall attempt to recover from the collision.

V_B1.2 Successive collisions (i.e., a collision during the recovery from a previous collision) may result in the robot shutting down all activity and declaring failure. The number collisions in a chain that the robot can tolerate varies.

The random exploration behavior coexists with all the other possible behaviors that we might define. In the absence of any obstacle or collision, the robot will potentially be performing some other functions which are defined by a subfamily. However, this family is *not* abstract; thus, if no other behaviors are specified the robot will move forward at full speed.

V_B1.3 In the absence of an obstacle or collision, the behavior of the robot may vary

C_B1.3 In the absence of an obstacle, collision, or any other specified behavior, the robot will move forward at maximum speed.

Door navigation: Maneuvering through a doorway is difficult for a mobile robot. Often, obstacle detection sensors provide little information about the environment; thus, doorways are often not seen as viable passageways. Furthermore, it is difficult for the robot to find doorways in the first place given the noisy sensor data it receives.

C_B2.1 The robot shall attempt to locate doors in its environment

C_B2.1a Once the robot has found what it believes to be a door, it shall not consider the sides of the door to be obstacles as the door is navigated. See [C_B1.1], [V_B1.1].

V_B2.1 The width of the door which can be navigated by the robot will vary according to the width of the robotic platform and the quality of the on-board sensors.

Environmental interpretation: This behavior allows the robot to signal when it encounters a particular object in the environment. That object or objects will be identified by a particular color.

C_B3.1 The robot will signal when it has detected an object in its environment of the desired color.

V_B3.1 The color of the object(s) to be detected will vary and may be configurable at run time.

6.3 The whole family

The real mobile robotics domain is significantly more complex than space allows us to present in this paper. For example, we have not discussed whether or not the robot can move objects in its environment (with a gripper, for example). Nevertheless, we can illustrate some interesting properties of the domain even with this limited example. Suppose that we had four mobile robots at our disposal:

Figure 7. The mobile robot family along the hardware dimension

- A custom robot made out of Lego pieces with two infra-red sensors in the front for obstacle detection, a front bumper, and tank-tread locomotion. We will call this one **LegoBot**.
- A Pioneer robot made by ActivMedia [1] which has an array of sonar sensors, a gripper, collision detection via motor stalled, and wheels for locomotion. We will call this one **Pioneer**.
- A Pioneer (see 2) with a color vision system. We will call this one **Pioneer w/ Vision**.
- A small “pickle” robot that can roll around, and jump over small obstacles, and that is equipped with a camera. We will call this one **Pickle**.

There are several ways of visualizing the mobile robot product family. First, we will examine the mobile robot family along the hardware dimension (Figure 7). Notice that family members can fall into one of four different categories. The robot may have only the basic capability, in which case it exists only for the family **Basic Platform**. This is the case for **LegoBot**. The robot may have either one or the other of the additional hardware capabilities specified by the **Enhanced Obstacle Detection** or **Environmental Vision**. Finally, the robot may possess both the additional capabilities of **Enhanced Obstacle Detection** and **Environmental Vision**; therefore, it lies in the intersection of those two subfamilies. This is only one slice of the system, however, and if we were to look at the mobile robot family along the behavioral dimension we would see a similar picture. A somewhat more effective means of viewing 2-dimensional product family is in a 2-dimensional grid as shown in Figure 8.

The representation is symmetrical in this case because of the one-to-one mapping between behavioral subfamilies and hardware subfamilies. The full mobile robotic domain, however, is not symmetrical. In the full domain, behaviors may be composed and combined to form a composite behavior. For example, we might envision a behavior which includes the door navigation, combined with a mapping function, a wall following behavior, and a high-level planner. The mapping and high-level planning behaviors will need to communicate with the lower level random exploration, door navigation, and wall following to direct the robot towards high-level goals. However, if the robot collides with an obstacle, then the lower level

Figure 8. A possible 2-dimensional view of the robot product-line

behavior will take over and recover from the collision. Thus the structuring of the behavioral dimension is much more complex and resembles Brooks’ subsumptive architecture [5]. Furthermore, defining the behaviors independent of the hardware allows us to focus on only the behaviors and their interactions (a significant problem in and of itself).

These combinations of behaviors might require several different sets in the hardware domain, which will have sub-families that define, for example, robots with grippers, robots with bumpers, robots that have radio communications devices, and so forth. Thus, it is generally not the case in the full domain that a behavior will require exactly one subset in the hardware dimension or that the behavior and hardware dimensions have the same structure. By defining the intersection of the hardware dimension with the behavioral dimension, we define which family members are viable and which are not.

The structuring technique presented results in the creation of more families within the domain than with a traditional approach. However, these sub-families are more cohesive and simpler than creating just one top level-family. We believe that this provides several benefits. First, the top-level family can now be much broader than was previously possible. Even the small mobile robotics example presented here is more complex, yet more elegantly represented than in [9]. Second, the overall family can be expanded and contracted by adding and subtracting sub-families. Finally, these techniques will allow a family to be more easily refactored as the definition of the family evolves over time.

One of the barriers to traditional approaches is that the whole organization must change to accommodate product-line oriented development. Many resources are required to develop the domain engineering support for the entire product line while at the same time continuing to produce products for existing customers. We hope that our approach allows an organization to start out with a high-level product family and reuse just a few key pieces between its major product areas. As the payoff from this reuse makes more organizations

resources available, the organization can then afford to make the family more rich (by refactoring and/or adding sub-families) and thus achieving more payoff from the effort.

Of course, these benefits do not come for free. The broader and more flexible view of product families allowed by our techniques will result in families which are more complex than traditional families. In addition, because of this broader view, it may be more difficult to determine what constitutes a viable family under our approach. Nevertheless, we feel that these techniques hold promise and may serve to advance the frontiers of product-line engineering.

7 Conclusions and future work

In this paper, we have reflected on some current issues with product-line engineering and presented our initial attempts towards extending the product family approach to better address these issues as well as support what we call *n-dimensional* and *hierarchical* product families.

We examined problems that we and others have had in developing product families. We concluded that the difficulties in modeling near-commonalities and variability dependencies stemmed more from the structure of the domain itself and the lack of an ability to adequately capture that structure.

Further, we presented a simple structuring technique based on set theory. This allows the analyst to capture the structure of the domain and not be biased by implementation or design concerns. This approach allows thinking about *n-dimensional* and *hierarchical* product families and encourages many different views of the system.

We presented a limited example of the approach, based in the mobile robotics domain. Due to space constraints, the domain had to be simplified immensely. Nevertheless, even the simplified domain that was presented illustrates the points of our approach.

In the future, we intend to continue developing this approach to product-line engineering. We are already working on several industrial-sized case examples with good initial results. We look forward to reporting the outcome of that work, as well as presenting a more detailed view of this approach. Furthermore, in the future, we intend to provide a more detailed description of the formal foundations of this approach and how it could be leveraged if the family members were described in a formal specification language.

Acknowledgements

The authors wish to thank Steven P. Miller at the Rockwell-Collins Advanced Technology Center for his thoughtful comments and encouragement with regards to this work as well as his insights on how the structuring techniques presented here might be applied to the systems manufactured by Rockwell International.

References

- [1] Activmedia robotics website. Makers of the Pioneer robot. <http://www.activrobots.com/>.
- [2] Mark A. Ardis and David M. Weiss. Defining families: The commonality analysis. In *Nineteenth International Conference on Software Engineering (ICSE'97)*, pages 649–650, 1997.
- [3] D. Batory and S. O'Mally. The design and implementation of hierarchical software systems with reusable components. *ACM Transactions on Software Engineering and Methodology*, 1(4):355–398, October 1992.
- [4] J.L. Bentley. Programming pearls: Little languages. *Communications of the ACM*, 29(8):711–721, August 1986.
- [5] Rodney A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14–23, March 1986.
- [6] Lisa Brownsword and Paul Clements. A case study in successful product line development. Technical Report CMU/SEI-96-TR-016, Software Engineering Institute, Carnegie-Mellon University, October 1996.
- [7] G. Campbell, J O'Connor, C. Mansour, and J. Turner-Harris. Reuse in command and control systems. *IEEE Software*, 11(5):70–79, September 1994.
- [8] Grady H. Jr. Campbell, Stuart R. Faulk, and David M. Weiss. Introduction to synthesis. Technical Report INTRO-SYNTHESIS-PROCESS-90019-N, Software Productivity Consortium, Herdon, VA, 1990.
- [9] Debra M. Erickson. Structuring formal requirements specifications for reuse: A mobile robotics case study. Masters Project, University of Minnesota, April 2000.
- [10] H. Gomma. Reusable software requirements and architectures for families of systems. *Journal of Systems and Software*, 25(3):189–202, August 1995.
- [11] Juha Kuusela and Juha Savolainen. Requirements engineering for product families. In *Proceedings of the Twenty-Second International Conference on Software Engineering (ICSE'00)*, pages 60–68, June 2000.
- [12] W. Lam. Creating reusable architectures: Initial experience report. *ACM SIGSOFT Software Engineering Notes*, 22(4):39–43, 1997.
- [13] W. Lam, J.A. McDermid, and A.J. Vickers. Ten steps towards systematic requirements reuse. *Requirements Engineering*, 2(2):120–113, 1997.
- [14] Robyn R. Lutz. Extending the product family approach to support safe reuse. *Journal of Systems and Software*, 53:207–217, 2000.
- [15] J. Neighbors. The draco approach to constructing software from reusable components. *IEEE Transactions on Software Engineering*, 10(5):564–574, 1984.
- [16] D.L. Parnas. On the design and development of program families. *IEEE Transactions on Software Engineering*, 2(1):1–9, March 1976.
- [17] R. Prieto-Diaz. Domain analysis: An introduction. *ACM SIGSOFT Software Engineering Notes*, 15(2):47–54, 1990.
- [18] David M. Weiss. Defining families: The commonality analysis. Technical report, Lucent Technologies Bell Laboratories, 1000 E. Warrenville Rd, Naperville, IL 60566, 1997.
- [19] David M. Weiss and Chi Tau Robert Lai. *Software Product Line Engineering: A Family-Based Software Development Process*. Addison-Wesley, 1999.