Proceedings of IDETC/CIE 2005
ASME 2005 International Design Engineering Technical Conferences
& Computers and Information in Engineering Conference
September 24-28, 2005, Long Beach, California, USA

DETC2005-85522

FEATURE RECOGNITION USING COMBINED CONVEX AND MAXIMAL VOLUME DECOMPOSITIONS

Eric Wang
Creative Design and
Intelligent Tutoring Systems
(CREDITS) Research Center
Sungkyunkwan University
Suwon, South Korea 440-746
wang@me.skku.ac.kr

Yong Se Kim
Creative Design and
Intelligent Tutoring Systems
(CREDITS) Research Center
Sungkyunkwan University
Suwon, South Korea 440-746
yskim@skku.edu

Yoonhwan Woo
Institute of Mechanical Technology
Sungkyunkwan University
Suwon, South Korea 440-746
ywoophd@hanafos.com

ABSTRACT

Next generation process planning systems should be capable of dealing with industrial demands of versatility, flexibility, and agility for product manufacturing. Development of process planning system is heavily dependent on feature recognition, but presently there is no satisfactory feature recognition system relying on a single method. In this paper, we describe a hybrid feature recognition method for machining features that combines three feature recognition technologies: graph-based, convex volume decomposition, and maximal volume decomposition. Based on an evaluation of the strengths and weaknesses of these methods, we integrate them in a sequential workflow, such that each method recognizes features according to its strengths, and successively simplifies the part model for the following methods. We identify two anomalous cases arising from the application of maximal volume decomposition, and discuss their cure by introducing limiting halfspaces. All recognized features are combined into a unified hierarchical feature representation, which captures feature interaction information, including geometry-based machining precedence relations.

KEYWORDS

Feature recognition, Negative Feature Decomposition, Maximal Volume Decomposition

1 INTRODUCTION

Next generation process planning systems should be capable of dealing with industrial demands of versatility, flexibility, and agility for product manufacturing. Feature recognition is a key technology to link design information and manufacturing information, and the development of process planning and cost evaluation systems heavily depends on it. Feature technology research aims at determining alternative component representations that form a suitable basis for a wide range set of activities throughout a product's life cycle.

Over the past decade, the technology for machining feature recognition has steadily advanced, and commercial feature recognition systems have appeared in the marketplace. However, current feature recognition systems are not powerful enough to accommodate industrial demands. These systems can interpret and process simple features, but proper handling of interacting features has been a major problem.

2 BACKGROUND

Successful feature recognition methods can be broadly characterized by their approaches: graph-based, hint-based, convex decomposition, and volume decomposition-recombination [7]. Graph-based approaches are most efficient at recognizing canonical instances of feature patterns, but generally cannot handle interacting features that do not match any feature patterns. Convex decomposition and volume decomposition methods are more robust in handling feature interactions, but are computationally intensive. Hint-based methods use various heuristics and rules to avoid the computational cost of exhaustive volume decompositions. No single method is best in all domains.

There have been attempts to develop hybrid feature recognition algorithms [7]. Laakko and Mäntylä [4] combine a graph-matching method with a rule- and constraint-based system. Gao and Shah's Minimal Condition Sub-Graph method [1] combines graph-matching and hint-based methods.

2.1 Convex volume decomposition

Alternating Sum of Volumes with Partitioning (ASVP) is a convex volume decomposition using convex hull, set difference, and cutting operations [2]. It organizes the boundary faces of a part in an outside-in hierarchy, while associating volumetric components with these faces. By systematically combining components according to the hierarchical structure of the decomposition and the face dependency information obtained during decomposition, the ASVP decomposition is converted into the form feature decomposition (FFD). For machining applications, positive-to-negative conversion is applied to all positive form features, which converts the FFD into a Negative Feature Decomposition (NFD) consisting of a positive base component and negative removal volumes. These negative removal volumes are classified as machining features according to their original face information [8]. This method is able to recognize interacting features.

The NFD approach generates rich feature information, including a hierarchical organization of the features, and feature interactions, including geometry-based machining precedence relations. However, its geometric domain restriction of planar faces, and cylindrical faces that interact with planar faces along circular edges, is a limitation.

2.2 Maximal Volume Decomposition

Maximal Volume Decomposition (MVD) is a volumetric decomposition-recombination method. It decomposes the delta volume of a part into volumetric cells by extending the halfspaces induced from faces of the part, then recombines cells into large and simple convex sub-volumes called maximal volumes [9]. Each maximal volume is bounded by halfspaces induced by faces of the part, and is convex. The delta volume is obtained as the set difference between the part and its stock material, where the stock material may be specified by the user, or generated automatically from the bounding box of the part.

It has been shown that most maximal volumes correspond to machining features. MVD can recognize intersecting features. However, it suffers from combinatorial complexity as the number of halfspaces in the part increases.

3 HYBRID FEATURE RECOGNITION SYSTEM

A key consideration in the development of a hybrid feature recognition system is to combine the different methods effectively so that they mutually complement each others' capabilities, without compromising performance and completeness. Our insight is that each individual feature recognition method can be characterized along an efficiency-versus-richness axis, representing a tradeoff between being fast but limited, or being thorough and robust. We propose a sequential workflow, in increasing complexity of the feature recognition methods, in which each method handles what it's good at, and successively simplifies the part model for the following methods. This necessitates a post-processing step to combine all methods' recognized features into a unified feature representation.

The major drawback of the NFD method is its geometric restriction. Due to its reliance on convex hull generation, it can handle machined parts having planar and cylindrical faces only. Conical, spherical, and toroidal faces (hereafter called CST faces) are outside NFD's geometric domain. Meanwhile, MVD can handle all analytic surfaces, including CST faces, but its cell generation approach incurs combinatorial complexity as the part complexity increases.

We propose a hybrid feature recognition system that combines a graph-based method (GRP), Maximal Volume Decomposition (MVD), and Negative Feature Decomposition (NFD), in exactly that order, as shown in Figure 1. GRP is described in Section 3.1. An extension to MVD called Selective MVD is described in Section 3.2. NFD is discussed in Section 3.3.

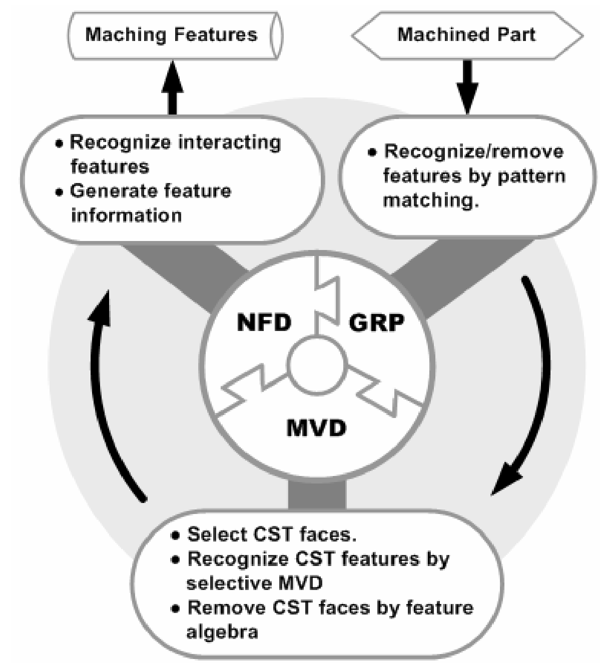

Figure 1. Hybrid feature recognition system

3.1 Graph matching method

The graph matching method (GRP) has been very popular in the feature recognition research community [3][5]. In this approach, the boundary representation of a part is represented as a face-edge graph, whose nodes denote faces and whose arcs denote edges. Features are defined as sub-graphs, and the part model is searched for occurrences of each feature sub-graph. This approach is very fast at recognizing non-interacting features, on the order of dozens or hundreds of features per second. However, when features interact, the topological and geometric patterns of the features are lost. Numerous techniques have been proposed to recover the missing topological elements, but with only limited success. Generally, the GRP methods cannot reliably handle interacting features.

We apply the GRP method first, primarily to simplify the part model before applying NFD. This obtains the set of non-interacting features. For each set of faces recognized as a feature, the corresponding feature volume is instantiated, and is removed from the part using Boolean operations, obtaining the filtered part with all non-interacting features removed. For the example of Part A shown in Figure 2, GRP recognizes two blind-hole features.

As a separate pre-processing step, we identify blends as cylindrical or toroidal faces with straight or circular edges that are tangent-continuous along two non-adjacent edges. Non-planar faces that are adjacent *only* to other blends are also identified as blending interaction patches, which we assume were generated automatically as a byproduct of the blending operation. These blends and blending interaction patches are removed from the part by unblending, which obtains the filtered and unblended part, shown in B in Figure 2. This is the input to the MVD method.

3.2 Selective Maximal Volume Decomposition

Maximal volume decomposition (MVD) can handle any analytic surfaces, including conical, spherical, and toroidal (CST) faces. However, its cell-based approach may suffer from combinatorial explosion. Selective MVD addresses the combinatorial issue by restricting MVD to a set of faces selected interactively by the user [9].

3.2.1 Recognition of features with CST faces

The selective MVD method has been further developed into a capability to recognize all volumes with CST faces, as follows.

1. Select all CST faces, and all faces adjacent to CST faces, in the delta volume, automatically.
2. Apply selective MVD using the selected faces to obtain sub-volumes:
   (a) Extend the selected faces and their adjacent faces that have concave edges.
   (b) Intersect the extended faces with the delta volume, obtaining cells.
   (c) Compose adjacent cells into sub-volumes if they satisfy the conditions for maximal volumes. (Some sub-volumes may fail these conditions.)

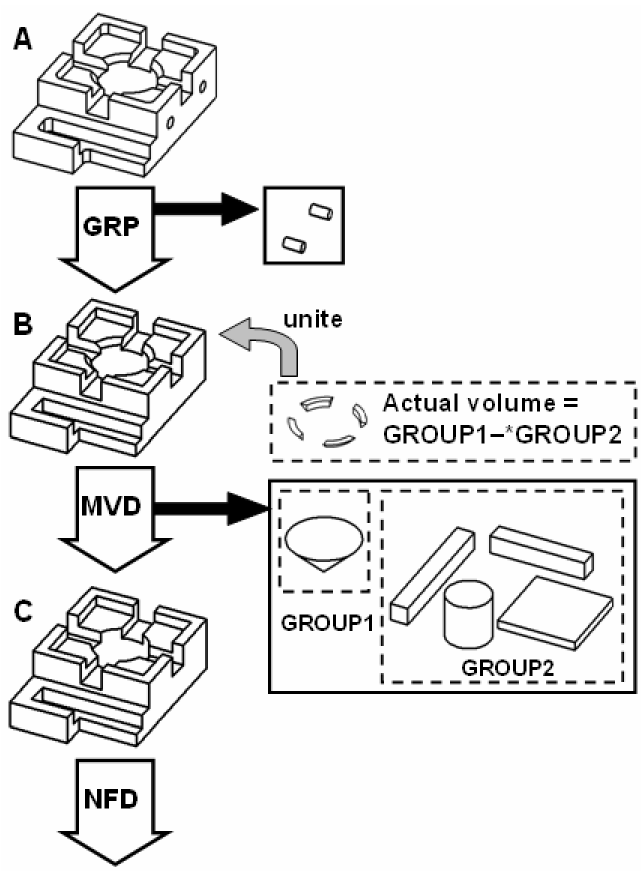

Figure 2. Hybrid feature recognition for Part A

Then the sub-volumes containing CST faces of the delta volume are the features with CST faces. One restriction in selecting CST faces is that only faces with concave edges are selected, since MVD requires concave edges.

3.2.2 Removal of CST faces

Given the sub-volumes generated by the selected MVD, a method has been devised to remove the CST faces from the part, using feature algebra.

1. Classify the sub-volumes based on face containment.
   (a) GROUP1: Sub-volumes that contain the CST faces.
   (b) GROUP2: Sub-volumes that have non-empty intersection with any sub-volume in GROUP1.
   (c) GROUP3: Sub-volumes that do not belong to either of the first two groups.
2. Generate the actual volume of each feature in GROUP1. For a feature F in GROUP1 that interacts with a set of features S_i in GROUP2, the actual volume of feature F is obtained as the set difference between F and the union of all S_i.
3. Unite the actual volume with the part.

For the filtered and unblended part B shown in Figure 2, the sub-volume with a conical face is in GROUP1, and it interacts with all four sub-volumes in GROUP2. Applying feature algebra obtains the actual volume shown in Figure 2. Uniting this with the filtered and unblended part obtains the non-CST part C shown in Figure 2, which has no CST faces.

3.3 Recognition of interacting features by NFD

NFD uses a recursive volumetric decomposition algorithm. This gives it a strong capability to handle feature interactions. It also generates rich feature information automatically, including outside-in volumetric hierarchy, face dependency, feature interactions, and geometry-based machining precedence relations. Alternative feature representations are generated by growing features through fictitious internal faces, and by aggregating features with common base faces.

Conversely, many non-interacting features could cause intermediate volumetric components to be highly interconnected, which may require remedial partitioning (cutting) operations during convex decomposition. Cutting operations are less desirable because they lose much of the rich feature information. We apply the GRP method first to simplify the part model, which can drastically reduce the need for cutting operations during NFD.

As the NFD algorithm relies on convex hull operations, it requires polyhedral input. The NFD method has been extended to handle blends and cylindrical faces within a broad geometric domain by defining conversion methods to and from a polyhedral abstraction. No such conversions have been defined for CST faces, so these cannot be handled by the NFD approach. Hence, we apply MVD first, to remove those features having CST faces.

3.4 Unification of feature representations

Each of the three methods outputs its recognized feature volumes separately. As the NFD provides the richest feature interaction information, we adopt it as the primary feature representation. Features recognized by GRP and MVD are then added to the NFD.

NFD automatically calculates face dependency relations [2] during convex decomposition. Given a GRP or MVD feature volume F, we obtain the equivalent face dependency relations as follows:

1. For each face *f* of F, find every original face *g* in the input part that has non-empty intersection with *f*, using 2D face intersection tests.
2. If any *g* are found, then *f* is an original face. For each *g*, look up every original face *j* in features J_j that are already in the NFD, such that *j* has face dependency to *g*. This involves trivial pointer traversals only, using *g*'s face pointer information stored in the NFD. For each *j*, add a new face dependency from *f* to *j*.
3. If no *g* is found, then *f* is a fictitious face. In this case, traverse the NFD hierarchy to find every face *j* that has non-empty intersection with *f*. Here, we can exploit the face pointer information to reduce the number of faces to be tested.
   - (a) For every face *a* adjacent to *f* in feature F, if *a* is an original face with dependency to a face *g* in the input part, then we look up every original face *j* with dependency to *g*, as described in step 2, and then test every face *k* adjacent to *j* in features J_j.
   - (b) If none of the faces adjacent to *f* are original, then we simply test every face *k* in the NFD.
   - (c) For each *k* that has non-empty intersection with *f*, add a new face dependency from *f* to *k*.

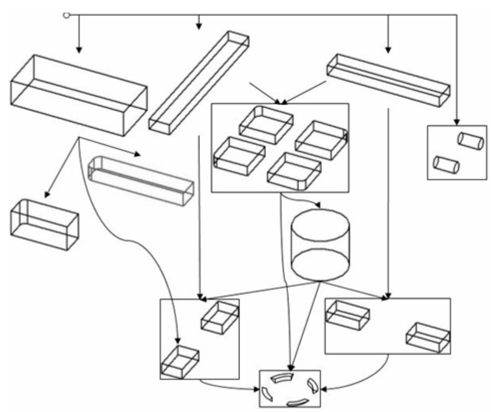

(a) Machining precedence relations

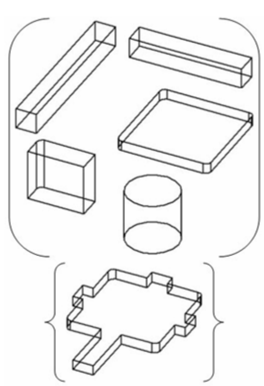

(b) Alternative features from growing and aggregating

Figure 3. Unified feature representation for Part A

After adding each new face dependency, we insert the new feature F into the NFD as a grandchild of existing feature J. Note that the NFD hierarchy is an "alternating sum of volumes". A quasi-disjoint union of two features of the same sign, e.g. both negative, is represented as a grandchild relation with a null intermediate node. That is, (J $\cup^*$ F) is represented as J $-^*$ ($\varnothing$ $-^*$ F).

Original face dependency relations are many-to-many in the general case. One face of the input part could be classified as original in two or more features. Conversely, one feature face could correspond to two or more disjoint original faces. Hence, a new feature F could have face dependency to two or more existing features Jj. In this case, F is added as a grandchild of all such features Jj. Thus, the unified feature representation becomes a (directed acyclic) graph structure.

The unified feature representation for Part A in Figure 2 is shown in Figure 3(a, b). This is a simplified visual representation of the NFD hierarchical structure which emphasizes the machining precedence relations, and illustrates the graph structure of the feature interactions. The actual volume recognized by MVD has been attached as descendent features of (i) the cylindrical hole, (ii) the four pockets, and (iii) the four slots.

4 ANOMALIES CAUSED BY MAXIMAL VOLUME GROWING, AND THEIR CURE

The Selective Maximal Volume Decomposition method of Section 3.2.1 can be thought of as a growing operation, since MVD sub-volumes are maximally grown up to halfspaces induced from faces of the part. This growing effect can result in anomalous feature output. We identify two anomalous cases, and describe their cure.

4.1 <u>Anomaly Type 1: Disjoint Features Merge</u>

Consider the case of a simple block solid with two small horizontal slots u and v that intersect with a vertical slot s, as shown in Figure 4(a–d). For simplicity, only 2D top views are shown. The bottom faces of slots u and v, shown as the blue and red rectangles, are coplanar, while the side faces of u and v, shown as horizontal lines in blue and red, respectively, are parallel, but not necessarily coplanar.

We distinguish four cases according to the relative offsets of these side faces: (a) neither pair coplanar, with both of u's side faces contained within the negative halfspaces induced from v's faces, (b) neither pair coplanar, with the four side faces interleaved, (c) one pair coplanar, and (d) both pairs coplanar. The result of MVD growing for each case is shown in the right column. In (a) and (c), the cell generated from red slot v is grown up to slot s's left side face, but no further. In contrast, the cell obtained from blue slot u is not blocked by s's right side face, so it grows through slot v to the boundary of the starting workpiece. In (b), both slots are grown up to slot s's side faces. For cases (a), (b), and (c), the two slots u and v grow to two distinct sub-volumes, preserving the intrinsic feature information of slots u and v.

In (d), the cells obtained from slots u and v are not limited by slot s's side faces, so they grow through each other. Hence, they are merged into a single slot sub-volume. This case loses some intrinsic feature information that was available in the original part, namely that the two slot features u and v were originally disjoint. We denote this case as anomaly type 1.

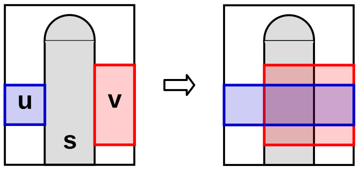

(a) {<, >}: Slots u and v have no pairs of coplanar side faces.

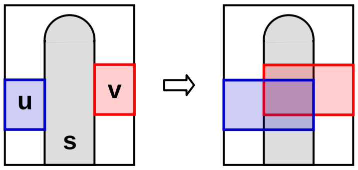

(b) {<, <}: Slots u and v have no pairs of coplanar side faces.

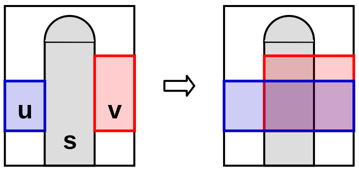

(c) { =, ≠}: Slots u and v have one pair of coplanar side faces.

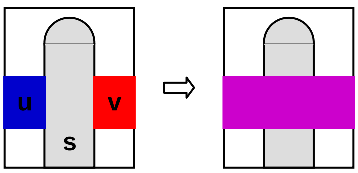

(d) { =, =}: Slots u and v have two pairs of coplanar side faces.

Figure 4. Maximal volumes for interacting slots

From a broad perspective, our role in feature recognition is to faithfully capture and interpret all features and feature interactions. It follows that all intrinsic feature information is potentially significant. Any changes to the intrinsic face information of the part could affect the feature recognition result. Especially, a feature recognition method's own internal algorithms should not introduce any unnecessary changes to this information.

4.2 Anomaly Type 2: CST Face Reappears

The feature algebra algorithm of Section 3.2.2 is intended to remove CST faces. Let F denote a feature having a CST face. Then F will be in GROUP1. When F's actual volume is united with the part in step 3, all of F's CST faces will vanish from the part. F's actual volume is obtained in step 2 as the set difference of F and all GROUP2 features that interact with F. But suppose some GROUP2 feature S has its own CST face. Then S's CST face could appear in F's actual volume in step 2. If it does, then this CST face would also appear in the part after step 3. Then the feature algebra algorithm has failed to remove all CST faces. This is anomaly type 2.

We note that whenever the MVD growing operation is halted by a halfspace induced from a face of the part, it introduces a new copy of that face, or a subset thereof, into the grown sub-volume. The limiting face could be any face of the part, even a CST face. Thus, a GROUP2 feature, even if it originates from a polyhedral cell, will have a CST face whenever it interacts with another feature G such that it grows up to a CST face in G. The feature algebra algorithm makes no guarantee that GROUP2 features are all polyhedral, so it cannot guarantee that it removes all CST faces. It is therefore incomplete.

4.3 Strategy for Curing Anomalies

Both anomalies arise from the intrinsic nature of MVD's cellular recombination, i.e. its growing operation. Hence, one possible strategy is to prevent the anomalies by revising MVD's growing strategy. We propose instead to accept MVD as an established method, without any major modification. This is justified on the grounds that MVD's capabilities and strengths have been extensively analyzed, and are well-known, and it is desirable to retain this body of work as much as possible, so as to leverage its benefits. Making a radical change to MVD's core algorithms could have unforeseen consequences that invalidate much of MVD's established strengths.

We therefore adopt the strategy of detecting and curing the anomalies as they occur. We will show that the anomalies can be detected and cured efficiently by introducing limiting halfspaces, so this strategy is effective, and is expected to have small or negligible runtime cost.

4.4 Curing Anomaly Type 1 using Limiting Halfspace

Part T in Figure 5 has a conical hole feature interacting with two small slot features, such that the slots' bottom faces, and both pairs of side faces, are coplanar. Note that Part T is equivalent to Figure 4(d). Hence, it also exhibits anomaly type 1. MVD generates a conical volume F in GROUP1, and a single maximal volume S1 in GROUP2, in which all three pairs of coplanar faces have been merged. Applying feature algebra obtains a non-CST part with a single large slot feature, as shown in Figure 5, upper right.

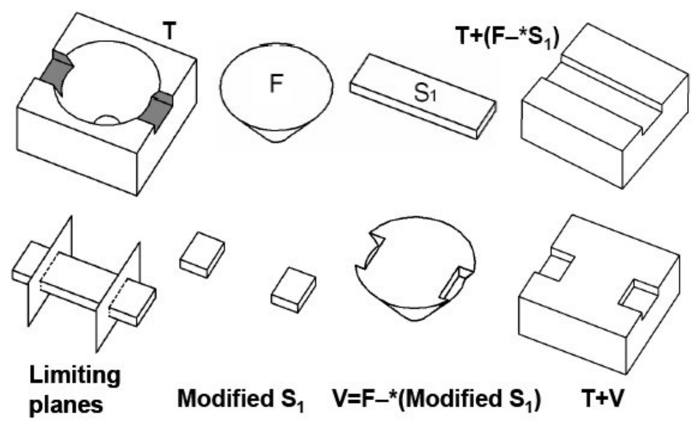

Figure 5. Curing anomaly type 1

We detect anomaly type 1 by counting original face pointers [2] from the faces of each MVD volume to the faces of the input part. Three faces of S1 are detected as having two original faces pointers each, originating from two disjoint sets of three connected original faces in T, which were the two small slot features in T. We cure this anomaly by introducing limiting planar halfspaces. For each set K of connected original faces in T, we introduce a limiting planar halfspace P within the volume of feature F, such that all faces in set K are in the positive halfspace of P, and all original faces in S1 not in set K are in the negative halfspace of P. The two limiting planar halfspaces for S1 are shown in Figure 5, lower left.

The result of feature algebra introduces two new planar faces into the non-CST part, which are classified as fictitious faces. Hence, the non-CST part contains two slot features, while preserving the intrinsic information that these two slots are disjoint.

4.5 Curing Anomaly Type 2 using Limiting Halfspace

Anomaly type 2 is illustrated for Part T in Figure 6. The conical feature F is in GROUP1. The slot feature in Part T, with three planar faces, interacts with F, so it is in GROUP2. MVD grows the slot through F to the conical face of F, obtaining MVD sub-volume S1, which has a CST face obtained by copying a subset of F's conical face. Applying feature algebra would cause S1's CST face to remain in the "non-CST" part, which is not desirable.

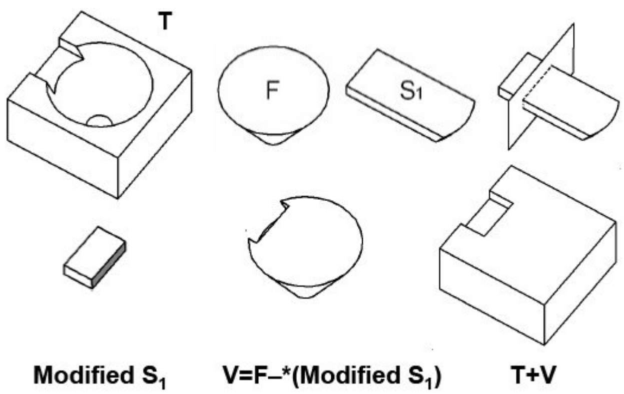

Figure 6. Curing anomaly type 2

To cure anomaly type 2, we consider a CST face to be an original face, i.e. originating from the part model, when it occurs in a GROUP1 feature, but as a fictitious face when it occurs in a GROUP2 feature. A fictitious face is not necessary in representing the shape of the feature, and can be replaced by any equivalent representation. We eliminate the fictitious CST face by introducing a planar halfspace anywhere within the volume of F, obtaining a modified S1, shown in Figure 6, lower left. Conceptually, this planar halfspace limits the extent of the maximal volume of S1. The limiting halfspace is verified using intersection tests to ensure that the modified S1 does not interfere with the part. Applying feature algebra using the modified S1 introduces a new planar face in the part, which is classified as a fictitious face.

5 APPLICATION TO REAL-WORLD PART

A change arm, shown in Figure 7, is a real-world part used to transport machining tools from a tool magazine to a milling machine. It rotates around a central attachment. The tool shaft is engaged in the end portion, which is primarily cylindrical to accommodate the tool shaft. Each end portion has a conical surface, which mates with a ring-shaped protrusion in a sleeve around the tool. This conical surface is the key surface that provides vertical support during transport.

Previously, feature recognition using graph matching + NFD could not handle the conical faces, so feature recognition was completed on a simplified change arm model with these conical faces removed. In the current hybrid feature recognition method, the conical faces are removed by MVD, which allows NFD to proceed as before. Thus, the hybrid method can fully handle the original change arm model. The recognized features for this part are shown in Figure 8.

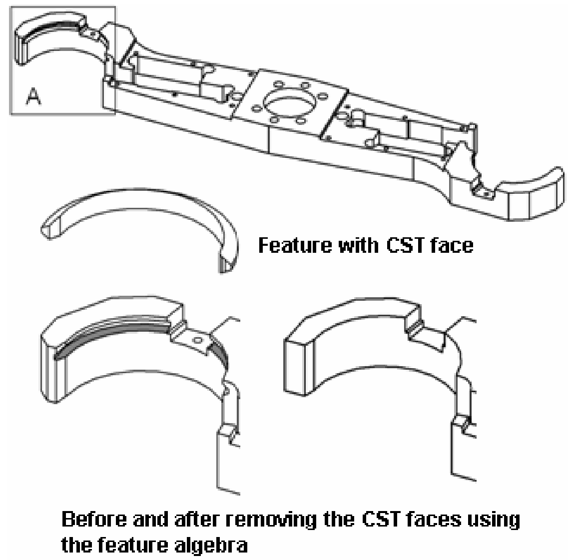

Figure 7. A change arm part

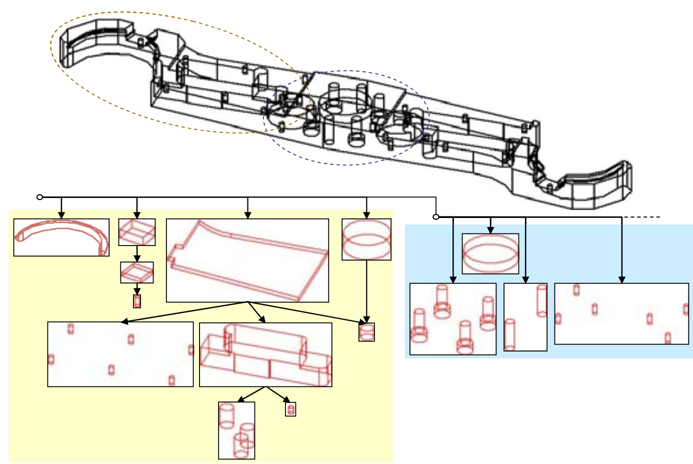

Figure 8. Recognized features for change arm part

6 CONCLUDING REMARKS

A hybrid feature recognition method for machining has been presented in this paper, which combines three distinct feature recognition methods in a sequential workflow. GRP is efficient, and tends to simplify the volumetric connectivity of the resulting part. Selective MVD removes features with CST faces, while addressing MVD's combinatorial complexity. NFD handles interacting features, and generates rich feature information. Anomalies caused by MVD's maximal volume growing have been identified, and are cured by introducing limiting halfspaces.

Each of these feature recognition methods has previously been implemented separately. The integrated hybrid system is currently being developed as part of a next-generation feature recognition and process planning server.

ACKNOWLEDGEMENT

This research was supported by the Korean Ministry of Science & Technology through the Creative Research Initiative Program.

REFERENCES

[1] Gao, S., and Shah, J., "Automatic Recognition of Interacting Features Based on MCSG", *Computer-Aided Design*, **30**(9), 727–739, 1998.

[2] Kim, Y. S., "Volumetric Feature Recognition Using Convex Decomposition", in Shah, J., Mäntylä, M., and Nau, D. S., Eds., *Advances in Feature-Based Manufacturing*, Elsevier, Chapter 3, 39–63, 1994.

[3] Kumara, S. R. T., Kao, C. Y., Gallagher, M. G., and Kasturi, R., "3-D Interacting Manufacturing Feature Recognition", *1994 CIRP Annals*, **43**(1), 133–136, 1994.

[4] Laakko, T., Incremental Feature Modeling: Methodology for Integrating Features and Solid Models, PhD Thesis, Helsinki University of Technology, 1993.

[5] Owondunni, O., Mladenov, D., and Hinduja, S., "Extensible Classification of Design and Manufacturing Features", *2002 CIRP Annals*, **51**(1), 133–166, 2002.

[6] Sakurai, H., Dave, P., "Volume Decomposition and Feature Recognition, Part II: Curved Objects", *Computer-Aided Design*, **27**(11), 833–843, 1995.

[7] Shah, J., Anderson, D., Kim, Y. S., and Joshi, S., "A Disclosure on Geometric Feature Recognition from CAD Models", *Journal of Computing and Information Science in Engineering*, **1**(1), 41–51, 2001.

[8] Wang, E., Kim, Y. S., Lee, C. S., and Rho, H. M., "Feature-based Machining Precedence Reasoning and Sequence Planning", *Proc. ASME Computers in Engineering Conf.*, Atlanta, Sep. 1998.

[9] Woo, Y., and Lee, S. H., "Destructive Modeling by Volume Decomposition and Its Applications", *Proc. ASME Computers in Engineering Conf.*, Chicago, Sep. 2003.