

Software Engineering Process Model — A Case Study

Gary J. Nutt

Department of Computer Science, CB 430

University of Colorado

Boulder, CO 80309-0430

nutt@cs.colorado.edu

Abstract

Recently a faction of software engineering researchers has focused their attention on studying the process by which software is produced, stimulating interest in models to specify, design, and implement software. A significant part of the practicing software industry must produce software that conforms to a documentation standard (military standard 2167A) for software products; it is intended to ensure that delivered software meets the documentation requirements. This paper is a case study of how a government software contractor might use models to define a process for designing and implementing a software product that complies with the documentation requirements. The intent of the paper is to apply business process modeling technology to the software engineering domain, thus exploring strengths and weaknesses of our evolving models of group collaboration. The case study illustrates an alternative way to design, analyze, and track software processes. It also attempts to illustrate how the model might “break down” as the basis of an enactment model if it were to be used to coordinate the work of a large number of software developers.

Keywords: Process models, workflow, software engineering

1 Introduction

Models have been used to describe work processes by decomposing the procedure into a set of discrete steps, then showing how information flows among them. The purpose of the process is arbitrary, and has traditionally varied from manufacturing processes to

information management strategies. Variants of these process models have been heavily-used for modeling specific behavioral aspects of organizations: PERT charts describe how manufacturing and engineering processes can be organized to build a product. Queuing networks model workflow through a system of service providers in terms of service and interarrival times. Flowcharts describe how a sequential program should execute to transform information.

We are interested in applying workflow modeling techniques to describe the process of creating software products. There is considerable research activity based on the idea that one can represent the steps in producing software as an algorithmic process [9, 10]. An important component of this research effort tends to cast *software process models* as programs in high level languages (e.g., see [12]). However, due to the importance of maturity models in software organizations [8], much of the software process work focuses on this aspect of process modeling rather than the earlier algorithmic specifications of the process.

The software crisis has long been known to be a critical technical problem, and has even been identified as a grand challenge problem. Very large software systems tend to be extremely difficult to manage for a number of reasons: it is difficult to create precise requirements for desired systems. There are few good ways to partition designs and to integrate the resulting modules. There are few good tools for managing the work. Organizations tend to dramatically underestimate the complexity of the target system (workers build far more complex units than the project managers can fathom). Despite years of research in the area, there is no proven methodology for software engineering.

The federal government contracts substantial software projects to independent contractors. Historically, the government has often been disappointed by the quality and nature of software that is delivered in response to a contractual commitment. Over the years, this has caused many government agencies to attempt to provide guidance regarding the process

Permission to make digital/hard copies of all or part of this material for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee.

COOCS 95 Milpitas CA USA © 1995 ACM 0-89791-706-5/95/08..\$3.50

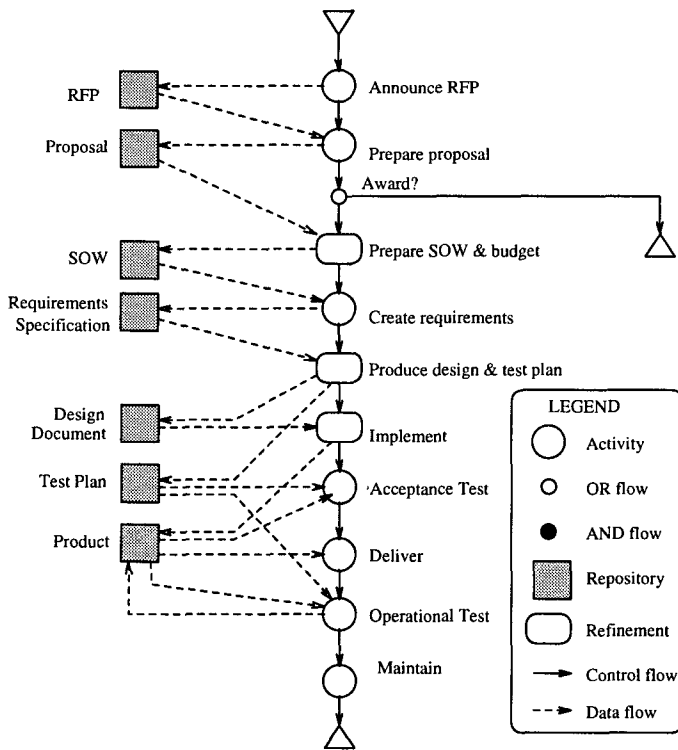


Figure 1: The Waterfall Model (Idealized 2167A Process)

that is used to produce the software (without actually prescribing an operational software methodology) so that they can have some assurance that software behaves as intended. The military specification 2167A is one standard that agencies may use to constrain the way that software is developed under a government contract [13].

The 2167A specification requires that the contractor use some software engineering environment, but avoids specifying which one to use. Any organization that intends to meet the specification requirements will need to employ some relatively confined process to be able to satisfy the requirement. As a practical matter, the standard implies some family of processes under which acceptable software can be developed.

This paper explores the use of process/workflow models to represent the software process for organizations that comply with the 2167A specification; the models are first used to consider *idealized* processes, then to consider various aspects of more realistic operation.

Under ordinary usage of the 2167A specification, there is a *granting agency* that desires to have some software product built, and one or more *bidding agencies* that could do the work. The granting agency announces a *request for proposals* (RFP) inviting bidding agencies to submit a formal *proposal* indicating

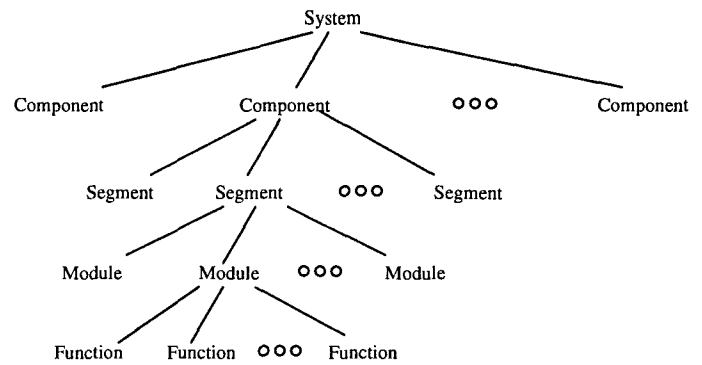


Figure 2: Hierarchical System Organization

how it would produce the software (see Figure 1). The granting agency selects one bidding agency, and negotiates an agreement based on the proposal. The result of successful negotiation is that the granting and bidding agencies agree on the *statement of work* (SOW) that specifies the precise deliverables and a budget. The agencies then commit the statement of work and budget to a contract with specific terms and conditions.

The bidding agency can then begin to do the work under the constraint that its work (demonstrably) conform with the documentation required by the 2167A standard. First, the deliverables listed in the SOW are converted to a *system requirement specification* that identifies explicit, testable requirements for each deliverable. The standard does not allow requirements that cannot be tested, otherwise it is not possible to determine if a deliverable satisfies the negotiated contract or not. For example, a requirement such as “the system shall be fast enough to avoid catastrophe” is not acceptable, while a requirement such as “the response to any command must be accomplished in 11 seconds” is acceptable.

The systems requirement document is used to drive the design; the design phase must produce a *system design document* and a *system test plan*. The system design document must specify:

External interfaces. The specification of all input operations and data for the system, and the corresponding result at all output ports for the system. Each individual requirement in the systems requirement specification must be satisfied by some aspect of the external interface.

Architecture. Describes how the system will accomplish the functionality required of the system. Functions at the external interface must correlate with specific parts of the architecture.

Subsystems. The architecture can be partitioned into subsystems. Required function implementations

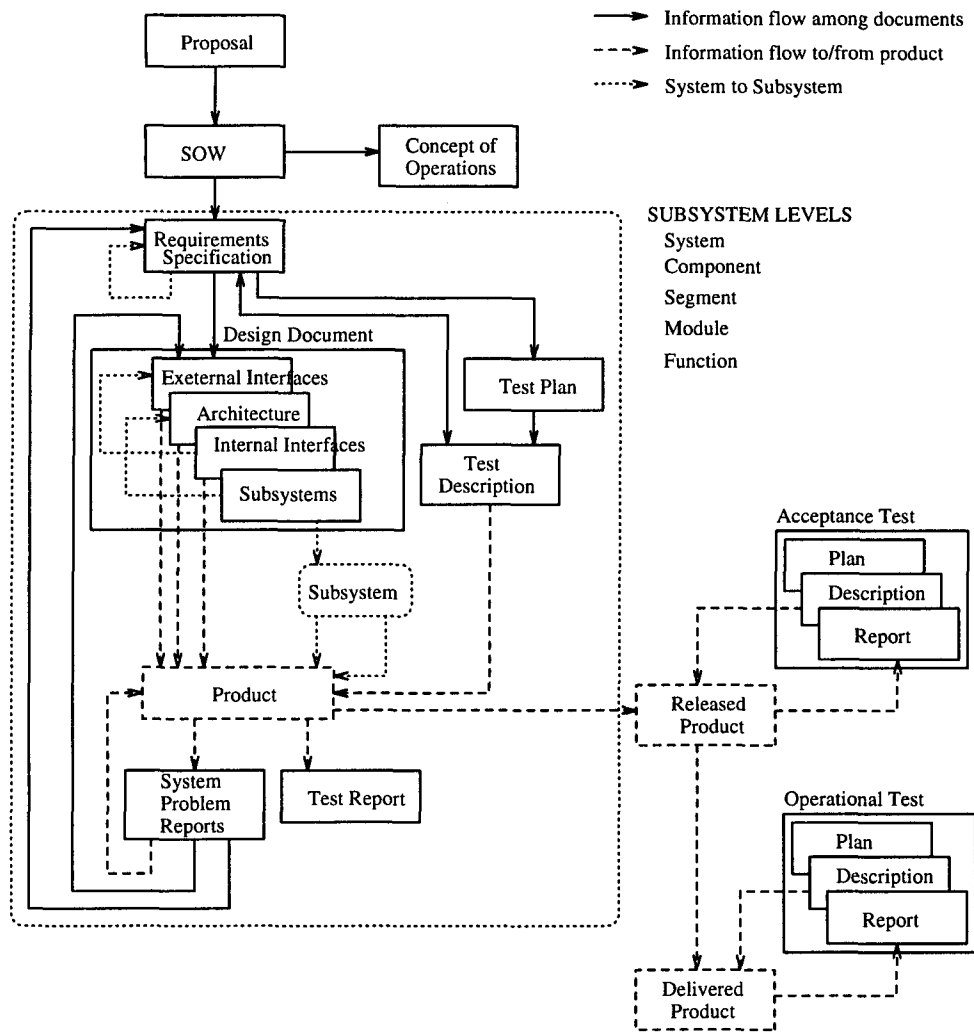


Figure 3: 2167A Document Dependencies

must be traceable to a subsystem.

Internal interfaces. The specification of the interfaces among subsystems. Each internal interface specifies an external interface for a subsystem.

The system test plan is prepared at the same time as the system design document, and specifies:

Test Plan. How each unit in the design will be tested to show that it implements specific requirements.

Test Description. The nature of the tests that can be applied to a design unit after it has been implemented to illustrate that it performs as required.

Test Report. A document that is produced when the test has been applied to the unit.

Part of the overall system test plan defines an *acceptance test plan* used at the time the product is delivered, and an *operational test plan* used at the time

the product is installed (often the acceptance test plan and the operational test plan are the same).

Because of the wide use of hierarchical decomposition in systems, the 2167A standard presumes the system will be organized into a hierarchy, e.g., with levels corresponding to the system, components, segments, modules, and functions (see Figure 2). Design documents and a test plan must be produced for every unit in the system, i.e., when a system is decomposed into a set of components, each component must have a design document and a test plan as described above for the system. In particular, subsystem requirements are derived from supersystem requirements, and the design document and test plan are derived from the subsystem requirements. Each function in the supersystem requirement must have a counterpart in the union of its decomposed subsystem requirements, and into the design document and test plan. Further, most organizations strongly discourage “vertical iteration”

in which subsystem design causes supersystem requirements and/or test plans to be changed; otherwise, individual groups cannot be delegated work without the danger of wasted effort, and at the top level this would imply that negotiated agreements with the granting agency might have to change.

Once the system has been completely designed it can be implemented. Implementation is accomplished by delegating work to different groups and allowing them to construct their assigned units of the end product. As each unit is completed, it is tested to ensure that it meets the test plan and hence the cascaded requirements. Integration is implicitly part of hierarchical development.

Ultimately, the full system is built and can be prepared for delivery to the granting agency. Acceptance tests are run at the time the system is released by the bidding agency and initially accepted by the granting agency; subsequent operational testing may take place when the product is installed in its target environment.

Figure 3 summarizes the documents, information flow among documents, and the relationship between the product and the documents. We use a dotted line to represent the hierarchical relationship of documents, i.e., the dotted line from requirement specification to requirement specification is intended to mean that the supersystem requirement specification is used to define the subsystem requirement specification. In particular, we use a dotted rectangle to surround part of the document flow, and a “recursive call” on this part of the process within the flow diagram.

2 A Process Model

Figure 1 is a very high-level view of the idealized (waterfall) process that a bidding agency might choose to employ. Subsequent figures refine the idealized model to provide more detail and alternate perspectives of the process an organization might use to produce software in a manner that satisfies the 2167A requirements.

The SOW is created through negotiation between the granting and bidding agencies (Figure 4). The proposal specifies the basic boundaries of the SOW under a specific budget; it is used as a starting point in negotiations between the two agencies. The response to the proposal is to add some functions to those originally proposed and to remove other proposed functions; the budget is adjusted to reflect the revised SOW. The model describes how a bidding agency might respond to the proposal; it must use the proposal, budget, and the response from the granting agency to create a specific SOW. In the figure, the activity entitled “Convert requirement to SOW” represents this work. The

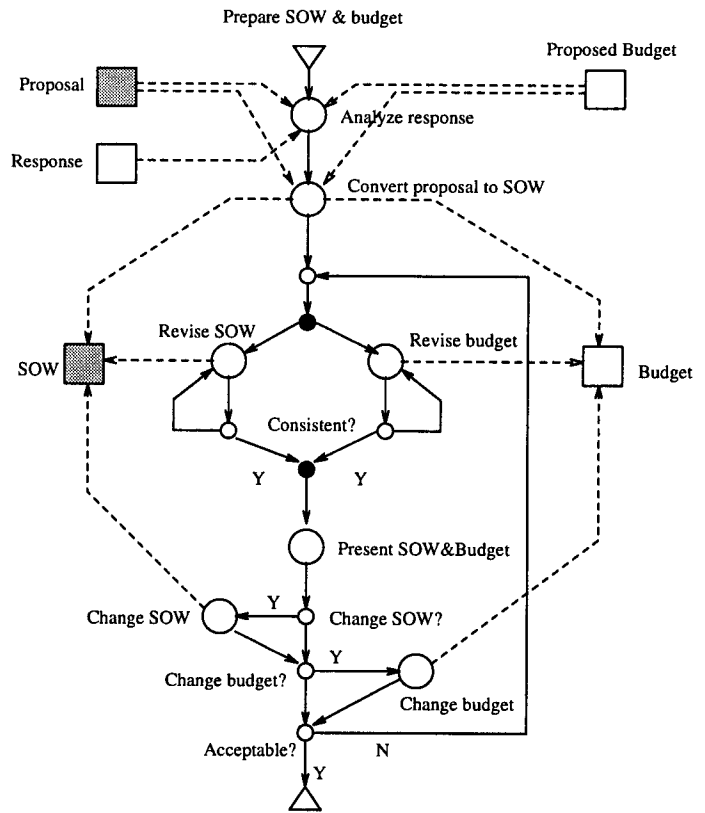


Figure 4: Preparing the Statement of Work (Idealized 2167A Process)

SOW and budget are then revised (logically in parallel); this revision process continues until the proposed SOW and budget are consistent. The revised SOW and budget are then presented to the granting agency for approval. The granting agency may accept the SOW and budget or may ask for either to be revised — see the figure. If revision is required, then the process loops back through the part of the process that ensures that the SOW and budget are consistent prior to presenting them to the granting agency again. This iteration continues until the SOW and budget have converged onto an acceptable plan. The bidding agency is then ready to create the system level requirement by translating the SOW into requirements (we do not refine this part of the process).

The Design Activity. Once the system requirements specification has been completed, the system-level design and test plans can be completed. Figure 5 describes more details of this process. This macro step represents a substantial fraction of the work — sometimes *all* of the work on a contract. The model is intended to describe the recursive topdown design process where any system can be decomposed into a set of subsystems: the first step is to create the test plan and

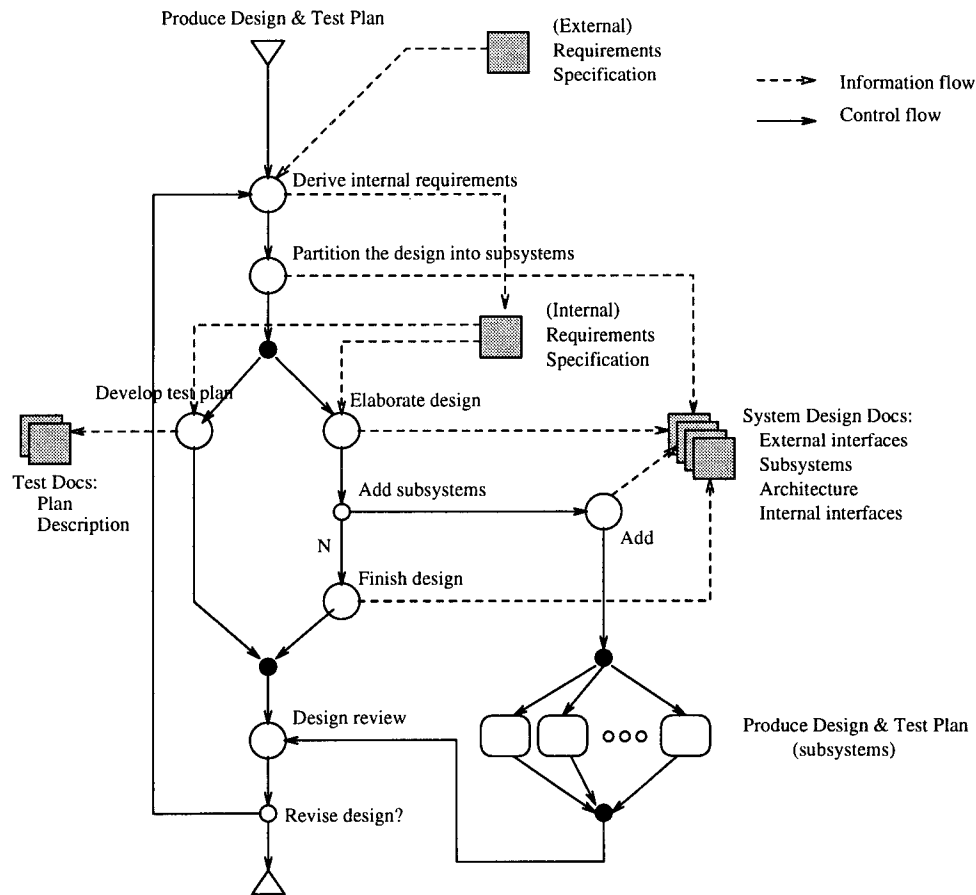


Figure 5: Produce the Design (Idealized 2167A Process)

design document for the overall system. Part of the work of determining the system design is to partition it into subsystems that can be designed independently, interacting across internal interfaces specified as part of the design at this level. Notice that the model illustrates that subsystem refinement is a *wholly nested* activity, meaning that the system design activity has a single entry and exit point with refinement resulting in path fan-out and fan-in contained within the internal details, but not visible at the activity interfaces. This property corresponds to the information in the waterfall model, i.e., all design is completed before any implementation is started; it also represents that subsystem design cannot effect supersystem requirements and tests plans. (However it rarely represents the way projects are actually accomplished.)

Bidding contractors approach the project using a process based on the waterfall model with the design elaboration represented by Figure 5. The maximum design depth is likely to be about five levels as suggested by the common subsystems level names (system, component, segment, module, and function), but the common usage has no implied limit or suggestion regarding the breadth. The standard and the model

both support arbitrary breadth and depth.

Figure 5 is a compact representation of the design step that represents a hierarchy of arbitrary breadth by eliding first generation subsystems, and of arbitrary depth by recursively referencing the “Produce design & test plan” submodel. When the model is used to represent any specific process instance, it results in a tree of design processes as suggested by Figure 6.

Idealized Implementation. Figure 7 represents a refinement of the “Implement” activity in Figure 1. The model again uses the ellipsis operator to represent depth refinement of the subsystem hierarchy, but models the breadth at any level using iteration. This follows since the design hierarchy results in a set of leaves, each of which should result in the creation of some unit of software — called a “function” in the figure. Functions are encoded (based on the function design), tested to see that they meet the interface requirements and the test plan; a (unit) test report is produced to verify that the function meets the requirements. (Ultimately, each unit test should be traceable back to a high level system requirement.) If the implementation does not pass the test, a system problem

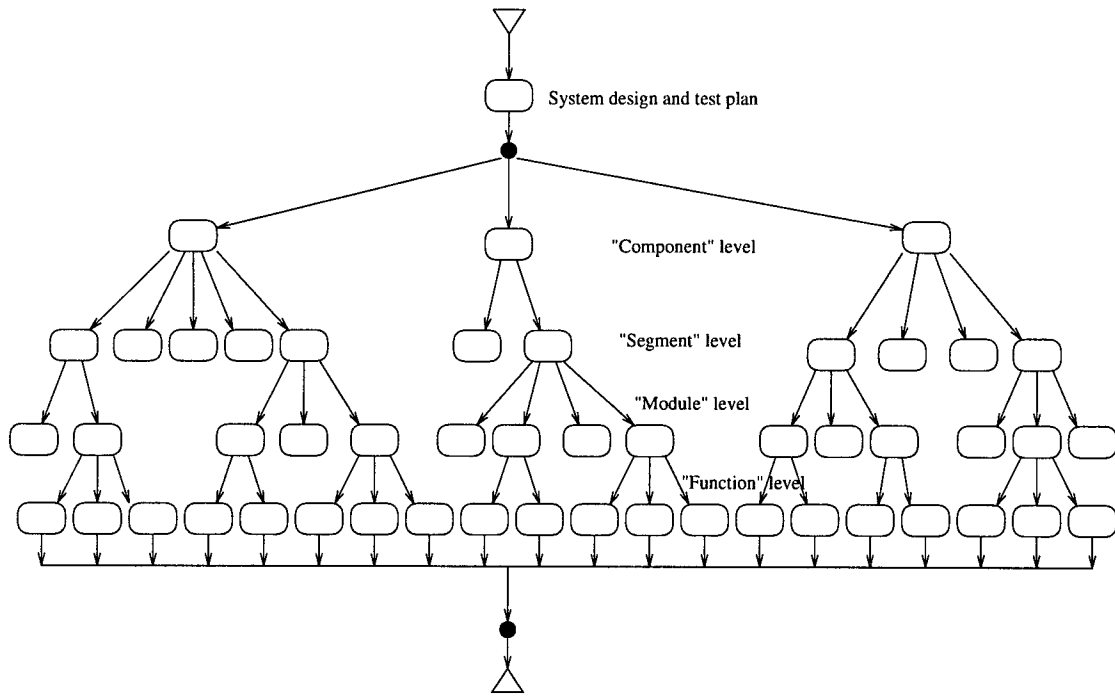


Figure 6: An Instance of the Design Hierarchy

report entry is created (not shown in the figure) and the function is refined and retested. Once a function has successfully passed the unit test, it is integrated into its encompassing supersystem — a subsystem in the design hierarchy. In the figure we have elided the various subsystem levels, representing subsystem testing and integration similar to that shown for function integration.

3 Practical Considerations

The model described in the previous section represents an idealized view of the software process. It is likely to be useful as a high-level description to explain the process to the granting agency or a new employee. However, it probably is not an accurate reflection of how the software would be created in the organization.

The first problem with the model is that the abstraction of the design and implementation hierarchies is too abstract for practical use. Before it really imparts the nature of a specific process instance, it is necessary to translate the subsystem elaboration into a model that represents the actual design hierarchy such as suggested by Figures 6. This suggests that if a model and system were to be useful to software organizations, the basic process could be defined using the system, then elaborated for each specific contract.

This representation also suggests a second problem; there are many tasks to perform in the process, but nothing to indicate how the work is delegated and accomplished. Our discussions with practicing software engineers is that the models represent the concept, but that they do not capture the essence of the enactment of activities. For example, if a “technical” person is responsible for the project, then various activities are deemphasized while others become the subject of intense attention; conversely, if a “manager” is responsible for the project, then different activities become the focus of activity — substantially different work is done in one case then in the other.

The third problem is the most serious problem (and it is easily recognizable from case studies of office systems): the idealized model represents the general idea, but it does not represent the way work really gets done. This is not due so much to change and exceptions as it is to the fact that the model simply represents the intent of the process rather than its mechanics.

Roles and Actors. The idealized model can be refined to specify more detail about the process, although the level of description provided in Section 2 may be adequate to address the various documents that are required by the 2167A specification. Each workflow activity represents some unit of work that must be addressed before the overall task is completed. In a model that describes the work of a group or orga-

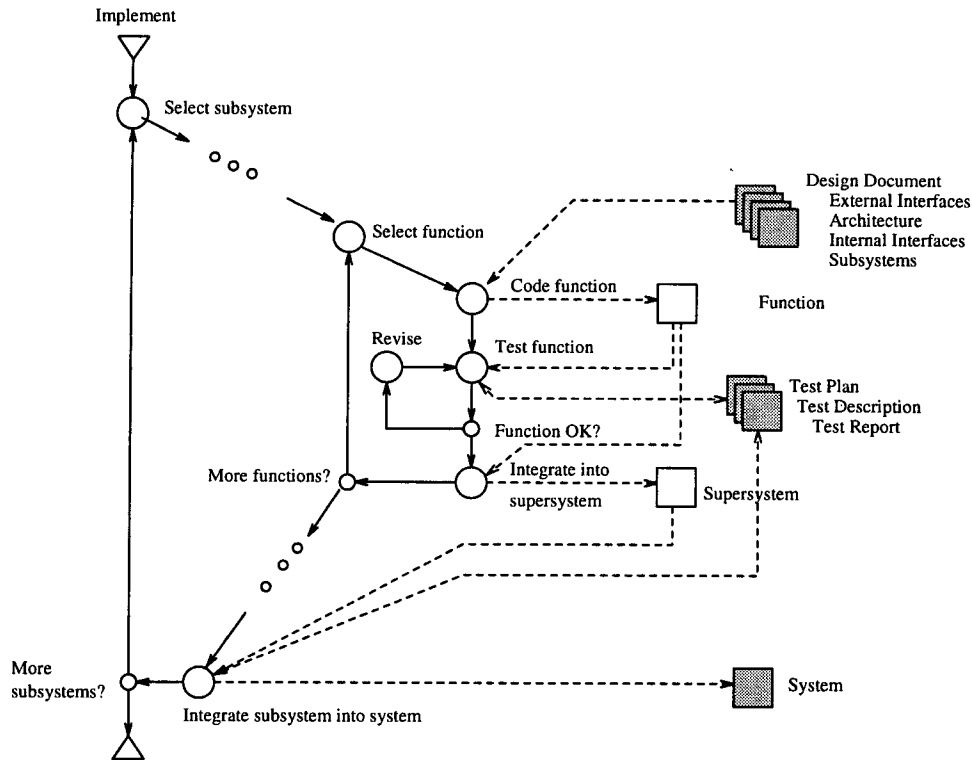


Figure 7: Implementation (Idealized 2167A Process)

nization, such as a software development organization, the model can be used to represent the assignment of work to various entities. A *role* designates a unit of work — one or more activities — that is to be fulfilled by an *actor*. Thus there is a mapping between activities and roles, e.g., system design activities are to be done by a “senior designer,” coding a function is to be done by a “junior programmer,” and subsystem testing is to be done by a “quality assurance engineer.” Conventionally this correspondence maps many activities to a single role, although there are arguments for allowing the mapping to be many-to-many. There is another mapping between roles and actors that specifies which specific person (or computer) is responsible for executing any particular role; of course there may be many actors assigned to any role, and an actor may have many different roles.

The nature of the organization begins to be evident in a role mapping, e.g., in some organizations it would be unusual for a programmer (or programming group) to be mapped into functions from two different modules managed by two different actors.

The actor mapping will also be determined by the group organization; while roles may be homogeneous across parts of the design hierarchy, work assignment (actor assignment to roles, and hence to activities) would correspond to management responsibility in a traditional hierarchical organization. If the organiza-

tion used a pool of programmers that could be assigned to any project, then the actor mapping would use an entirely different philosophy.

The model indicates that certain documents are to be prepared (depending on the design decomposition), in a particular order, e.g., the test plan is supposed to be derived from the requirements rather than from the design. There is a “super role” associated with managing the whole process; if the actor that fills this role is concerned with the quality of the documentation, then he or she will tend to closely manage the document preparation activities, but perhaps pay less attention to the nature of the activities that effect the actual design, e.g., the subsystem refinement at any given level. Further, some actors may follow the activity precedence religiously, while others may allow activities to be enacted in orders other than that specified by the model. (Of course out-of-order enactment can easily lead to violations of *intended precedence* in the model, e.g., the test plan development and design document should be concurrent and independent — if the test plan is designed after the design document, then it will tend to be derived from the design and the requirements rather than just from the requirements.)

The way that a role is fulfilled also relates to where it exists in the model. For example Figure 6 makes no distinction between the level of documentation related to system level design and function level design, while

in fact, functions may be implemented with very brief (or no) design document and test plan, relying on the supersystem to cover these aspects of documentation. (For example, an external policy could not necessarily distinguish between a part of the design hierarchy in which a module is a leaf and one which is further decomposed into functions.)

While there is a tendency to characterize these differences in actor behavior by the occupation of the actor (manager or technical person), it is also influenced by the nature of the group (e.g., is it market-driven or technology-driven), the size of the group (e.g., large groups need to have more rigorous rules and adherence to rules than do small groups).

Enactment versus Intent In this process *no* part of the implementation can proceed until *all* parts of the design are complete. In a realistic project, many aspects may not be known until far into the project; this will tend to cause gross inefficiencies in the organization — workers will remain idle waiting for their part of the design (which is complete) to be released until all other (related and unrelated) parts are all completed. Organizations rarely behave in this manner; instead, parts of the design are reviewed and released for implementation while other parts are still under design. Of course this strategy represents a different sort of gamble; if the design of an apparently unrelated part of the system suddenly becomes related (or even *causes the previously-reviewed and released design to change*, then the implementation work is wasted effort — we discuss this problem below). Thus Figure 8 might represent a more aggressive strategy for staging design and implementation (even though it is inconsistent with the higher level model). Various parts of the implementation are enabled to begin when their respective designs are complete. (In the figure we have used rectangles to replace parts of the design that are decomposed into hierarchical subsystems as is done in the 2167A document.)

This brings out an important point about using models as an aid for thinking and communicating about processes: high level models should be interpreted as an expression of general intent in the same sense that an abstract or executive overview summarizes a report; lower level models may be inconsistent with regard to formal modeling properties while being completely consistent with respect to the intent of the process.

We believe that the purpose of the 2167A specification is to force the bidding contractor to produce software products using a topdown methodology that conforms to the requirements. The contractor is expected to determine a set of requirements from the negotiated statement of work, to produce a test plan

that shows how to test each aspect of the requirements *independent of the design*, to create a design that satisfies the requirements and which can be tested, and to implement the design. The problem that bidding contractors have with the process is that it is very difficult to do using a topdown methodology for at least the following reasons:

1. Requirements are rarely complete enough to specify what is really expected from the product.
2. Complex requirements may be inconsistent.
3. Hierarchically-refined designs may uncover design errors at higher layers when lower layers are designed. For example it may not be possible to define a subsystem that meets the higher layered requirements.
4. Design problems may be discovered at implementation time.
5. Changes discovered in detailed design or implementation may impact the work assigned to other organizations; this will cause the other organization to be less efficient (e.g., to go over budget or over schedule), thus such changes will be met with rigorous resistance.
6. Persons responsible for high level contractor resources may not be able to be fully aware of design conflicts.

While these problems can become intractable in large software engineering organizations, it is clear that they are less related to the methodology than to management, organization, and to general intra-contractor politics. The granting organization simply requires that deliverable products be consistent with a negotiated contract. The problem arises due to the lack of understanding of the deliverable by both parties, and from the complexity of the deliverables.

Figure 9 represents a modification of the process that adds feedback links between the implementation and design phases of the waterfall model. When detailed design or implementation uncovers an argument for changing the design, then the appropriate “project manager” actor should determine how the project should handle the problem. The decision may be to ignore the problem by continuing with the original plan; this means that the resulting deliverable may be inefficient, or may fail some tests. The decision may be to change the design, implying that the impact of the change must be considered (e.g, how many other design units will have to change? What is the current investment in those designs and implementations? etc.) The process must then be rolled back and restarted.




Figure 8: Staged Design and Implementation

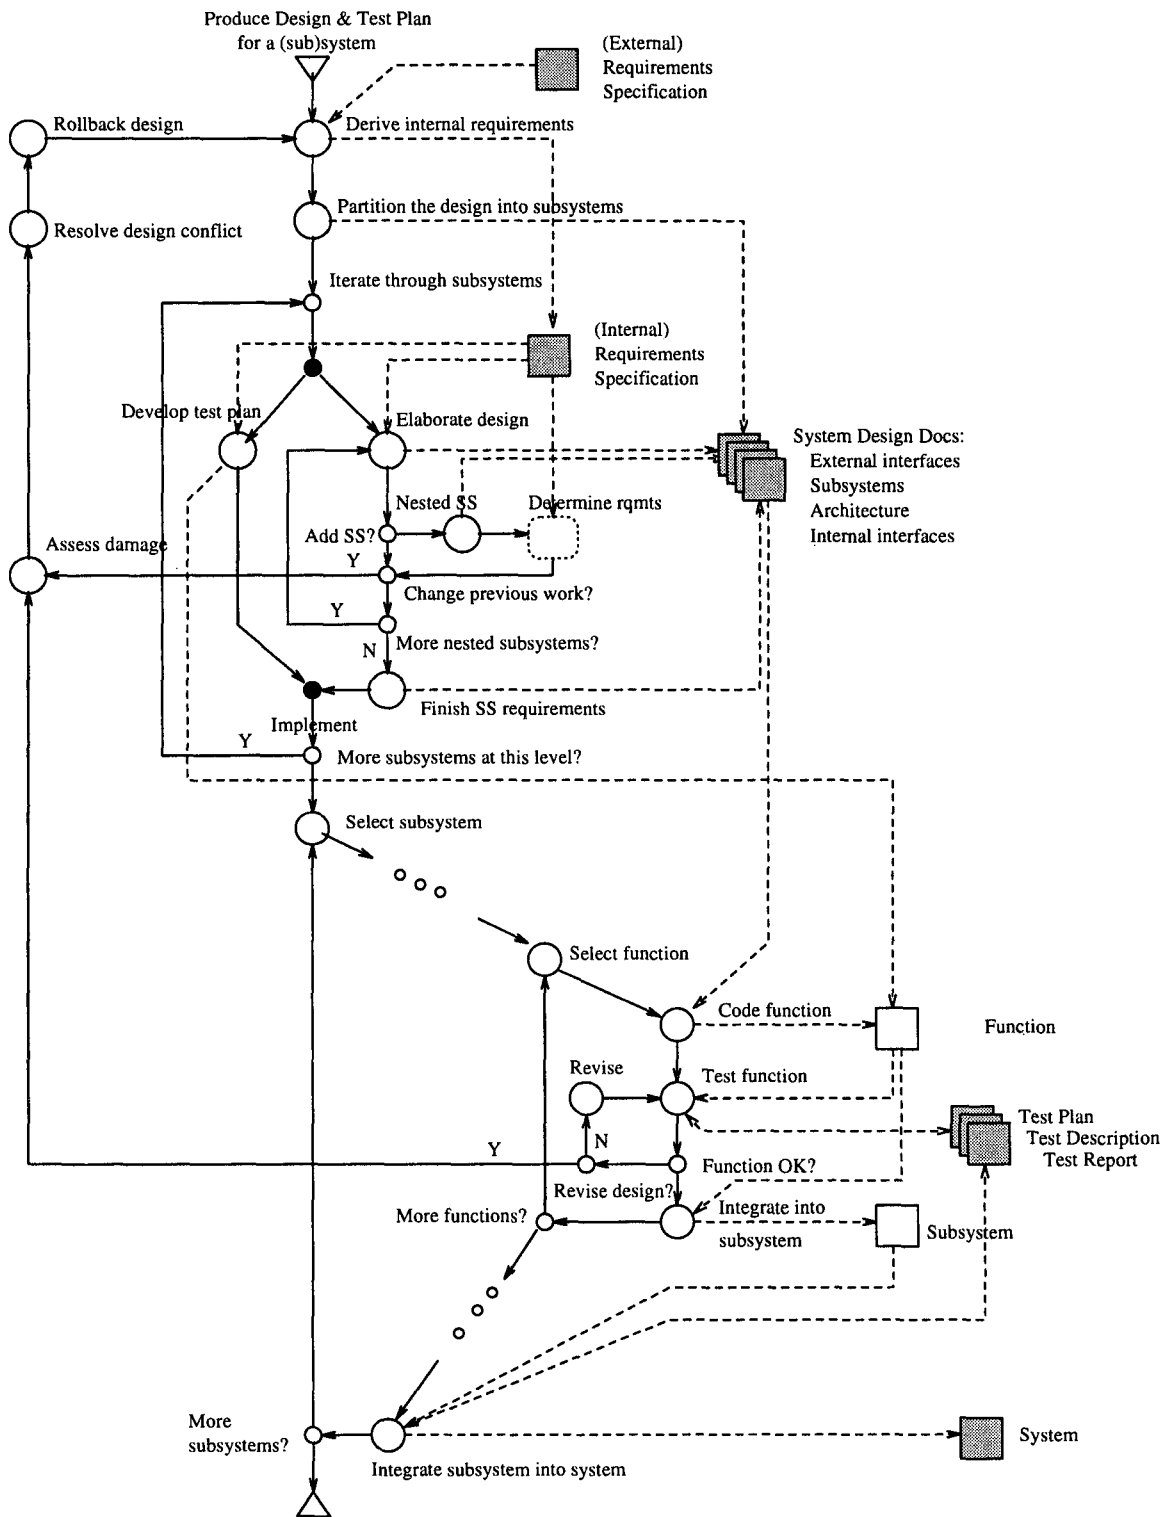


Figure 9: Exceptions in Design and Implementation

We only show design exceptions arising at two different places — during detailed design and during implementation. They could, of course, arise in other places, with corresponding edges leading to the “Assess damages” activity.

Process models are the messenger, not the problem. Our position is that the more information about the impact of changes that can be presented to decision makers involved in such decision the better. The process model is a language explicitly designed to convey that kind of information to a human.

4 Conclusion

We have provided a high level process model describing how different parts of an organization can interact to produce a software product along with documentation that conforms with the 2167A requirement. A more complete description of the model appears in [11]. Increasingly detailed models represent increasingly complex interactions among roles — groups — in the organization. The models we have illustrated are insufficient to “automate” the process, and would probably only be of limited use as an automated coordination model to stage work for the engineering organization. However, even at this level of detail they can serve several other important purposes when supported with a system:

Describing the Process. Large software projects frequently employ subcontractors and have staff turnover — new people arrive as the project progresses into detailed design and implementation, and staff depart after the product is released. In large projects, attrition is also normal, requiring that various staff members be replaced by new people. The model provides a high-level description of how the project is organized. When models of the form used in Figure 6 are generated, they provide a graphic description of the various levels and units in the design.

Software contractors continue to make extensive use of subcontracting organizations to handle parts of a contract. The 2167A standard requires that the contractor be able to ensure that the subcontractor complies with the specification. The descriptive model is a first step in satisfying this requirement.

Task Assignment. The model is a graphic basis for considering various strategies of organizing work teams and delegating work to them. Analysis tools enable a manager (at any particular level) to experiment with different strategies for assign-

ing work, particularly when used with animation facilities.

Tracking Tool. Auditing tools can also be added to the support system to track the status of the development work. If certain parts of the development begin to lag behind others, then the tracking facilities can be used in conjunction with task assignment facilities to adjust the number of workers focused on any particular part of the project. People responsible for specific units must sometimes make decisions about modifying a unit (recall Figure 9). The model provides a specific artifact to assist that person in making decisions regarding incorporating change, rolling back related development, etc.

Document Management. In a large project, document preparation is very difficult. The full document set must be internally consistent, and support function tracing from requirements to implementation (and the inverse). As designers create documents, they need to be able to navigate through the set of existing and pending documents to produce a document that is consistent with the rest of the system. While it is possible to design a database tool to maintain required relationships among documents, including generating skeletal documents when a design is initiated, it is difficult for a human user to logically absorb all of the relationships at once. The model provides a means by which the designer/implementer can navigate through the existing documents to understand the requirements and/or develop requirements for nested systems that are consistent and preserve traceability.

This case study has investigated the application of process models (ICN workflow models) to represent a software engineering process. We have not built a system to provide the support described above, although we have built several nontrivial workflow systems for description, modeling, analysis, and enactment. Our previous applications have addressed office applications, business processes, and distributed program organization, while this paper has focused on software processes. Nevertheless, the static model has proven to be useful in discussions with software engineers familiar with development under the 2167A guidelines. The primary benefit to these engineers has been in providing a simple, consistent, visual model that describes the process and the relationships among different components. The alternative view seems to present a new language of expression of the software process that experience software engineers know, but have difficulty describing to novice software engineers.

Our research group continues to study process models in an effort to make them more flexible for use in offices [2, 4, 5, 6, 7]. This case study has highlighted the need to be able to add a broad spectrum of different analysis tools to operate on a model, e.g., to manage requirement traces and to produce hierarchically consistent design and test documents. Like our experience with ICN-based enactment systems [3], this application continues to enforce the need for clear and complete mappings of activities to roles and to actors. Again, practical models identify the need for hierarchy in the model, but also point out problem with refining the model to include details that interact across subtrees of the hierarchy (e.g., Figure 8 illustrates a set of relationships between elements of the model in Figure 6 that is not evident in more abstract views of the same part of the process).

The study has emphasized the value of complementary model views for representing parallel activity (we also encountered this in our related work on parallel program models [1]): when one is focusing on understanding the process in general, descriptions analogous to Figure 5 are useful; however, once the model has been derived and it is to be used to study a particular instance of the process, views analogous to Figure 6 are more useful.

Finally, partially as a result of this case study we have refined our recent thinking regarding general models of processes so that we realize the value of descriptive models for such applications. Such models may lack precision in terms of strict analysis and enactment, but — when combined with animation — provide a rich means for conveying ideas about processes among humans.

Acknowledgement

The author was supported by Grant Number IRI-9307619 from the National Science Foundation. Most of the explanation for how bidding contractors operate under the 2167A guidelines came from Scott Brandt; he also reviewed and critiqued various interim models. The author and Skip Ellis continue to work on refinements to ICNs; that work is reflected in these models and ideas for systems.

References

[1] Adam L. Beguelin and Gary J. Nutt. Visual parallel programming and determinacy: A language specification, an analysis technique, and a programming tool. *Journal of Parallel and Distributed Computing*, 22(2):235–250, August 1994.

[2] Richard Blumenthal. *Representing Unstructured Work Flow Activities by Dynamically Exposing Contextual Information (tentative title)*. PhD thesis, University of Colorado, 1994. in preparation.

[3] Bull S. A. *Introduction to FlowPATH*, May 1992. Manual No. 44 A2 60XM.

[4] Clarence A. Ellis and Gary J. Nutt. The modeling and analysis of coordination systems. In *ACM 1992 Conference on Computer-Supported Cooperative Work*, 1992. workshop position paper.

[5] Clarence A. Ellis and Gary J. Nutt. Modeling and enactment of workflow systems. In *Advances in Petri Nets 93*, pages 1–16, June 1993. Invited paper.

[6] Clarence A. Ellis and Jacques Wainer. A conceptual model of groupware. In *ACM Conference on Computer Supported Cooperative Work*, 1994. To appear.

[7] Clarence A. Ellis and Jacques Wainer. Goal based models of collaboration. *Collaborative Computing*, 1:61–86, 1994.

[8] Watts S. Humphrey. Characterizing the software process: A maturity framework. *IEEE Software*, 5(2):73–79, March 1988.

[9] *Proceedings of the IEEE Software Process Workshop*. IEEE Computer Society Press, 1993.

[10] IEEE Software theme on Mature Processes, July 1993. Robert Lai, Guest Editor.

[11] Gary J. Nutt. Software engineering process model case study. Technical Report CU-CS-760-94, Department of Computer Science, University of Colorado, Boulder, December 1994. Abstracted version appears in COOCS 95.

[12] Leon Osterweil. Software processes are software too. In *Proceedings of the 9th International Conference on Software Engineering*, pages 2–13, Monterey, CA, 1987.

[13] U. S. Department of Defense. *Military Standard, Defense System Software Development*, 1988. DOD-STD-2167A.