

1.3 Foundational Concepts

Along with the preceding formulation of terminology, an appreciation of several general concepts (software development as a discipline, process and product quality, and the causes of product change) provides a foundation for understanding the approach being described.

The Evolution of Software Development as a Discipline

Every systematic discipline of practice has started as an opaque and individualized handcraft. As experience is gained, tools and techniques are developed that reduce routine and more mechanical aspects of the effort. Fabricated parts are created to further reduce the detail work required to create a finished product. At this point, it is still a fundamentally “manual” effort that relies on the skills and efforts of individuals to fashion more or less custom products. As understanding grows of an overall standard process for creating a product, machinery is developed for automation and mass production in which efforts begin to focus more on improving productivity and product quality. Improvements in mass production ultimately lead to mass customization for products that are individually tailored to specific needs and preferences, partially regaining some of the benefits of handcraft but with greater consistency and reduced labor.

Software development is following this same progression. It is now at the stage of being a tool- and parts-supported craft within a somewhat standardized process. This is encumbered by a mentality focused on building unique, statically-defined products that solve well-understood problems.

This craft-based approach has been augmented over time, with more disciplined practices, introduction of tools supporting more mechanical aspects of development (e.g., configuration management and builds, documentation, regression testing), and greater use of previously developed parts (i.e., commonly useful functionality).

Although the basic principles and practices of software engineering are well-established based on 50 years of experience, the form these take will continue to progressively change. These and other advances are a foundation for an industrial mindset that will

supplant systemically-flawed assumptions of traditional practice leading to the ability to mass produce customized products.

Separation of Concerns

Separation of concerns is the principle of organizing a complex matter into a constituent set of elements (i.e., of action or information). Each such element is conceived to have singular authority over a specified scope of responsibility and can be understood and realized in terms of its unique contribution to the matter as a whole.

This principle guides the structuring of, for example, a process into activities, an enterprise into programs, human activity into roles, and hardware and software functionality into discrete components. This is the basis for the concept of information hiding that Parnas conceived as the foundation for designing software for ease of change.

This principle motivates distinguishing between engineering and manufacturing as distinct elements of product development. Manufacturing is seen as the realization of products whereas engineering is seen as providing the means for effective and efficient manufacturing. A third element, management, is seen as the direction and coordination of engineering and manufacturing efforts within an enterprise.

The Need for an Engineering Discipline of Software

In early days of building software, when problems to be solved were simpler, code could be written using a simple notation to just mimic the steps a person would take for a solution. Code was written in discrete sections (e.g., blocks or subroutines) that provided sufficient computational separation of concerns. Such code did not have a complex structure and could be quickly understood by others. The translation from a human description of the problem to a machine-processable solution was simple. The existing way of working was affected only by the selective automating of discrete, previously manual tasks.

Software was and still is built by devising a human expression of intended behavior that is then translated into a notation that a computational device can follow to enact that behavior. The original conception of this required developers to define behavior in

the native language of computational hardware; subsequently, stylized notations (i.e., programming languages) were created to allow developers to describe behavior in a more human-readable form that could then be mechanically translated into a computational device's native notation.

In due course, as problems have become more complex and solutions more long-lived, code needed to have a more clearly defined structure for subsequent understanding of how it should be expected to behave and what to change if different behavior was needed. With obscure relationships among different parts of the code, changes in one part can indirectly affect other parts. Behavior of the code has become harder to predict and more likely to be different than expected. As a result, more time has to be spent trying to be sure the software works correctly. After all that, the software may still work differently than users had imagined or needed. As users' needs change over time, the code becomes still more complicated due to unplanned changes and yet more difficult to get correct and then change further. Of greater concern, the time from recognizing the need for a product and having that product available for use has grown increasingly long and even then often requires users to change how they prefer to work for the solution to work effectively.

In best cases, software is developed within an enterprise that has a specific focus and experience building particular types of software. Practitioners in those enterprises have the competence to build such software more efficiently and effectively than those who lack such competence. In this context, a product can often be built using a mix of new and previously developed components. A conventional approach makes no overt accommodation for this, leaving it to individual developers to recognize and exploit such opportunities, or not. This is a highly undisciplined and inefficient way of leveraging past work and misses the opportunity to lay a proper foundation for building and evolving software-based products.

With the growing complexity of problems and solutions, an engineering discipline is needed for building software-based products. Engineering is the systematic application of knowledge and expertise to achieve a solution to a problem, considering properties and tradeoffs among alternatives to determine a solution that best satisfies all aspects of

the problem. In building a software-based product, the problem to be solved is often initially not fully understood or properly communicated. As potential solutions are explored, understanding of the problem improves until alternatives can be sufficiently reduced and criteria for an acceptable solution determined.

Over time, as experience with similar problems increases, systematic methods arise for more clearly defining a problem at the start and reliably focusing on and then resolving among alternative viable solutions. Furthermore, the perception that two problems are similar is a basis both for expecting them to have similar solutions and equivalently for expecting that the solution to a past problem can be reasonably modified to be the solution to a similar problem. With similar problems, less effort is needed to determine and build customized solutions to each, with a focus on having to address only important differences among them. This works best if the developer understands how and why such problems are amenable to similar solutions and how differences in two problems lead to differences in their solutions.

Process and Product Quality

The “value” of a thing is its utility for a given purpose and its quality in serving that purpose. Quality is the degree to which the thing approximates the “ideal” form of such a thing for a given use. Quality is ultimately a subjective discernment but, for an enterprise to prosper, it benefits from having the means to continually improve the quality of its efforts. Quality improvement involves having measurable goals for its endeavors, mechanisms for determining how well those goals are being achieved, and actions to be taken when ways are found to improve both guidance for and associated practice of endeavors.

An enterprise is concerned with two sorts of quality: process quality and product quality. *Process quality* is the efficiency and effectiveness (equivalent to “productivity”) of an effort intended to achieve some envisioned result. *Product quality* is the degree to which a result (i.e., a product or service) approximates its potentially achievable best form for satisfying its purpose.

Process Quality

A defined process specifies how actions are organized to accomplish some purpose, comprised of a partially ordered set of activities and the practices to be used for each activity. Process quality represents the degree to which a properly defined and performed process is effective for reliably achieving an intended result.

The objective of a development process is to build a product whose capabilities satisfy some need. The goal of productivity is to reduce the delay between recognizing the need for a product and the consequent effective operational use of that product. This goal persists as needs change over the useful life of the product, requiring repeated modifications. Process quality for development of a software-based product (elaborated in section 2.0 as “developmental” quality) is evaluated in terms of four aggregate properties: feasibility, sustainability, conformability, and verifiability.

Quality Perspectives on Process Variation

The effectiveness of an effort differs depending on the suitability of the process used and the ability of practitioners to effectively perform that process. There are three perspectives on process variation that inform achieving a higher level of productivity in building a given level of product value:

- **Capability** – The range in productivity achievable with a given process, varying dependent on the complexity and uncertainty of a problem-solution
- **Maturity** – The degree to which a process is aligned with practitioner competence so as to enable predictably achieving a targeted level of productivity
- **Performance** – The productivity achieved for a given product value by practitioners following a given process

The objective for productivity is the high-maturity performance of an appropriate high-capability process. Process maturity is the alignment of an effective process with the competence of practitioners in performing that process and its associated practices. The capability of a process is the limit in productivity that practitioners can achieve in the performance of that process.

Productivity can be improved from several perspectives: (1) automating aspects of the process to reduce manual effort, unplanned changes, and product defects requiring rework, (2) improving process-practitioner alignment by revising the process and associated practices to better fit practitioner competence, (3) providing training and mentoring of practitioners in best use of the process and associated practices as defined, and (4) modifying the process or practices to streamline, automate, or eliminate activities resulting in a higher range of feasible capability.

Product Quality

Product quality represents the degree to which a product is suitable for its intended use. Product quality is dependent on the process with which the product is realized in that the process determines the criteria and attention given to achieving the various elements of product quality. Product quality can be improved either by increasing the effort spent on evaluating and achieving product quality goals or by enhancing the process with more effective means of addressing product quality. Product quality for a software-based product (elaborated in section 2.3 as product “behavioral” quality) is evaluated in terms of four aggregate properties: functionality, performance, dependability, and usability.

{note that “functionality” is defined here as a behavioral quality factor in that it will need to be traded against other factors (i.e., restricted in order to achieve e.g. security, safety, or performance goals).}

Expending greater effort on product quality can be viewed as reducing productivity unless the resulting improvement in product quality is recognized as requisite to product viability, whether initially or over its useful life. The challenge is to find a proper balance between acceptable quality and timely cost-effective deployment of a needed product. The process cost of improved product quality can be mitigated by improvements in the process that increase potential productivity, reducing the effort needed to achieve given product quality goals.

A more indirect dependence exists between product quality and the quality of the customer’s operational process: the effectiveness of the product depends on its

capabilities and use being a proper fit with the customer's process and practices. Another dependence is between development process quality and the quality of provider enterprise business practices (discussed in section 2.1) (e.g., legal/ financial, human resources, technical infrastructure) upon which development practices rely.

Directed Reviews

A directed review is a moderated interaction between an author of a product model content and author-designated reviewers to resolve author-specified concerns and identify additional issues related to the author-specified scope of the review. Additional reviewers may be given the opportunity to comment or identify issues in response to specific author-specified questions regarding given content. The author designates 2-4 reviewers as having relevant expertise or having responsibility for related content. Each review follows a directed approach¹ as follows:

- Specific content of the product model, and any relevant dependencies with other content, is the designated subject of a review.
- The characteristics of each reviewer involved in an element review.
- How the competence of each reviewer corresponds to the content to be reviewed.
- The questions that are posed for each reviewer to answer based on their particular competence relative to specific aspects of the reviewed content.
- Specific assertions that each reviewer makes about reviewed content.

Quality Reviews

Evaluating the degree to which process and product quality objectives are being met relies on ongoing timely reviews of results. Quality reviews follow a goal-question-metric, or similar, discipline. In this, the developer of material to be reviewed identifies appropriate reviewers and focuses each reviewer with questions appropriate to their

¹ Based on D.L. Parnas and D.M. Weiss, "Active Design Reviews: Principles and Practices", *Proceedings of the 8th international conference on Software Engineering*, August 1985, Pages 132–136.

specific competence related to particular content so as to focus them on issues of most significance, complexity, or uncertainty to increase confidence in that content. Questions to be considered should address adherence to developmental quality criteria, proper satisfaction of technical objectives including implications for achieving behavioral quality criteria, and consistency and completeness relative to all related specifications, realizations, and other related materials.

The Causes of Product Change

A product is built to provide capabilities that support the needs of customers in a targeted market. Although it is natural to think of a product as being the fixed realization of some associated solution to a specific problem, both the problem and its solution are susceptible to change. A product is better viewed as a changing solution to a changing problem.

Three factors lead to a product being changed: problem-solution uncertainty, product deficiencies, and problem-solution changes. A realistic approach to development must account for each of these factors for a product to be sustained over its expected useful life. These three are all interrelated in that each of them can arise even as either of the others are being addressed (e.g., a deficiency that is the result of an incorrect resolution of some uncertainty, an identified problem-solution change that changes how a problem-solution uncertainty should be resolved, or a problem-solution uncertainty that is not recognized until a problem-solution change occurs.)

Problem-solution uncertainty is a consequence of problem-solution complexity and the limits of human capability (or quoting Dijkstra: “On our inability to do much”). Product deficiencies express when the product is found to be flawed, having defects that degrade its ability to satisfy its envisioned purpose and capabilities, or when the product as released is acceptable but incomplete. Problem-solution changes are the result of changing customer needs, including changes in operational circumstances or enabling technology.

Problem-Solution Uncertainty

A product is a particular solution to an understanding of a problem. However, based on two of the precepts of software development (from section 1.1), uncertainty is a concern in two respects: is the actual problem properly understood and which of its many potential solutions is a best fit for the problem? The essential challenge, and substantial effort, of product development is resolving problem-solution uncertainty.

An effective solution requires a reasonably correct and complete understanding of the problem and the behavior that the product must exhibit to provide the customer with the capabilities needed. Certainty is an unobtainable ideal but uncertainty has many sources potentially subject to mitigation for improved confidence in the problem and a suitable solution. Uncertainty about the problem arises from unknowns (including excluded information due to security / proprietary / competitive concerns), tacit or overlooked assumptions, inherent complexity, miscommunications, ambiguity, and differing user or expert opinions about the problem or solution based on prior experience. These uncertainties can usually be satisfactorily resolved through customer-developer collaboration in exploring subject matter knowledge, customer objectives and practices, and circumstances related to operational context.

Uncertainties about the problem and its possible solution demands an iterative “learning” process of development through which understanding of the problem improves and may change and the nature and tradeoffs of alternative solutions are explored and comparatively evaluated. A premature or ill-founded resolution of uncertainties gives the illusion of problem-solution understanding but results in development of a deficient product.

Development of a product begins with a general conception of capabilities that the product should provide for its intended purpose. Such a conception lacks specificity in that many different products could in principle satisfy an insufficiently precise description of the problem or its solution. The process of creation is one of refining this conception by identifying and comparing alternatives in the problem and solution, considering issues of feasibility, practicality, constraints, convention, and preference imposed by the context of usage, to determine what form the product will take. The

effort involved in achieving a satisfactory product is dependent on the complexity of the actual problem and the viability of creating an appropriate solution to that problem.

Achieving a good solution is expedited with competence gained from solving similar problems but uncertainty can remain when alternative solutions fail to adequately resolve problem-specific conflicts among related aspects of product quality. Uncertainty about the solution concerns what combination of development choices will result in the right functionality with appropriate quality criteria, while satisfying any extrinsic constraints on the solution. Determining the best solution entails understanding the nature of the context in which the product operates and the interactions that result, how the solution addresses problem-imposed quality criteria and solution constraints, and how the computational platform affects product behavior with different solutions. Each of these aspects brings some degree of uncertainty that must be resolved to determine an acceptable solution.

A particular type of uncertainty relates to how understanding of the actual problem can change as development proceeds. An inaccurate understanding of the problem can result in resolving solution uncertainty that ends up as a deficiency in the product. Furthermore, the actual problem itself or recognition of possible solution alternatives can change even as the product is being developed, warranting changes in the solution to provide the customer with more effective capabilities.

Solution Deficiencies

A product is typically a satisfactory solution to the problem it addresses, seldom a perfect solution. The complexity of most software, and the challenges entailed in its development, means that it is going to have undiscovered defects. Furthermore, tradeoffs made among quality factors can lead to unforeseen emergent effects that could be improved with changes in the solution to achieve a closer fit to the problem. (quality factors are not absolute, every solution is an approximation to an envisioned ideal)

In addition, a sound practice of software development is to build a product that is a useful but incomplete solution to the full problem being addressed. The practice is to incrementally develop an effective solution to a partial problem, incrementally

addressing more aspects of the problem for an augmented solution in a series of product releases. The initial release provides a coherent set of capabilities but subsequent releases improve or add capabilities that the customer needed. Developing a product in this incremental way may result in the observable behavior of each release of the product purposefully differing and, in some cases, expand the market to which the product applies.

A customer's use of a product sometimes reveals that their endeavors would be improved with changes in product capabilities. Such changes may be needed for the customer to change how they perform an endeavor rather than simply to better support existing practices. This experience may be considered by the customer as being a deficiency in the product when it is actually a needed change in the problem or solution that the product has been properly built to realize.

Problem-Solution Changes

The suitability of a product can degrade over time unless the product is changed to accommodate changing user needs or to fit changing technology or changing operational circumstances. This mirrors the issues of problem-solution uncertainty in deciding how the product needs to change. Furthermore, such change can occur even as a product is being developed, with implications for reconsidering how aspects of uncertainty have already been resolved.

Recognizing the potential for change, building a product that is a fine solution to a properly understood problem is a good start but not the end. Changes can occur even as a product is being developed due to how uncertainties and tradeoffs in the problem or its solution are resolved or changes in the operational context in which the product is to be used. Over its expected useful life, changes in context will repeatedly occur in ways that will require changing the product so it continues to serve its intended purpose.

Modifying an existing product is substantially the same as dropping into the middle of initial development, recognizing that prior resolutions of some uncertainties are not easily modified. When a customer's needs or operational circumstances change, motivating reconsideration of the problem-solution for the product, the same

considerations of uncertainty and tradeoffs of the original development continue to apply. This will be an easier task if past uncertainties were resolved in expectation of these changes.

{relation to product deficiencies: changed perception of needs} In using a product, customers often discover improvements in how their enterprise operates that would benefit from changes in the product.

A particular concern in both creating and modifying a product is maintaining the ability to make further changes as customer needs and circumstances change. This requires judgement concerning the uncertainty regarding what types of changes are most likely. To be able to anticipate such changes, information needs to be provided, by problem and solution experts, as to what changes are more likely to be encountered during the expected life of the product.

Failing to identify likely potential changes is equivalent to assuming that nothing about a product will ever need to be changed. An informed but imperfect projection of potential change is generally more accurate than an (implicit) assumption of no change. Developing a product without considering likely future changes results in a product for which any change is unpredictably difficult and may be unnecessarily costly to make.

The degree of certainty concerning the nature of specific changes guides how much effort should be taken to prepare for those changes, recognizing that some potential changes may never occur or may take a form that differs from current expectations. Considering potential future change does not necessarily incur additional cost in that software developers often explore alternatives in order to resolve uncertainties and augment missing information (but then discard resulting information and insights that might save time when future changes do occur). More information about potential future changes serve to better inform developer explorations and allow them to better organize their work so as to localize likely changes.