

The Inevitable Pain of Software Development: Why There Is No Silver Bullet

Daniel M. Berry
Computer Science Department
University of Waterloo
Waterloo, Ontario N2L 3G1, Canada
dberry@uwaterloo.ca
10 August 2003

ABSTRACT

A variety of programming accidents, i.e., models, methods, artifacts, and tools, are examined to determine that each has a step that programmers find very painful. Consequently, they habitually avoid or postpone the step. This pain is generally where the programming accident meets requirements, the essence of software, and their relentless volatility. Hence, there is no silver bullet.

Categories and Subject Descriptors

D.2 [Software]: Software Engineering; D.2.0 [Software Engineering]: General; D.2.1 [SE]: Requirements/Specifications—*elicitation methods, methodologies, tools*; D.2.2 [SE]: Design Tools and Techniques—*evolutionary prototyping, structured programming, top-down programming*; D.2.3 [SE]: Coding Tools and Techniques—*structured programming, top-down programming*; D.2.4 [SE]: Software/Program Verification—*formal methods*; D.2.5 [SE]: Testing and Debugging—*code inspections and walk-throughs, tracing*; D.2.6 [SE]: Programming Environments—*integrated environments*; D.2.9 [SE]: Management—*life cycle, software configuration management, software process models*; D.2.10 [SE]: Design—*methodologies*

General Terms

Documentation, Economics, Human Factors, Management

Keywords

build-and-fix model, changes, configuration management, daily build, documentation, E-type systems, extreme programming, formal methods, inspection, open sourcing, pain, programming environments, programming methods, programming models, programming tools, rapid prototyping, regression testing, requirements engineering, silver bullet, software accident, software essence, structured programming, traceability, volatility, waterfall model

0 RATIONALE

The call for papers (<http://www.dsi.unive.it/~mont2002/topics.html>) for the 2002 Monterey Workshop on *Radical Innovations of Software and Systems Engineering in the Future* says of the object-oriented programming methods introduced in the last decade, “There is no proof and no evidence that software productivity has increased with the new methods.”¹ The call for papers argues that as a consequence of this and other reasons, there is an urgent need for new “post object-oriented” software engineering and programming techniques. However, there is no proof, evidence, or guarantee that any such new technique will increase productivity any more than object-oriented techniques have. Indeed, the past failures of any method to significantly increase productivity is the strongest predictor that any new technique will fare no better. In other words, what makes you, who designs new methods, think you can do better?

This paper tries to get to the root of why any given new programming technique has not improved productivity very much. This paper is, as the call for papers requires, an attempt “to analyze why some ideas were or were not successful” with the choice of “were not”.

1 INTRODUCTION

This paper is about building computer-based systems (CBS). Since the most flexible component of a CBS is its software, we often talk about developing its software, when in fact we are really developing the whole CBS. In this paper, “software” and “CBS” are used interchangeably.

This paper is based on personal observation. Sometimes, I describe an idea based solely on my own observations over the years. Such ideas carry no citation and have no formal or experimental basis. If your observations disagree, then please write your own rebuttal.

Sometimes, I give as a reason for doing or not doing something that should not be or should be done what amounts to a belief or feeling. This belief or feeling may be incorrect in the sense that it

¹ Actually, this claim is a bit too strong. There is a feeling that object orientation has improved programming a bit, and there are even some data [33], but it is clearly not the silver bullet that it was hoped, and even hyped, to be.

is not supported by the data. Whenever possible, I alert the reader of this mode by giving the belief-or-feeling-based sentence in italics.

2 PROGRAMMING THEN AND NOW

I learned to program in 1965. I can remember the first large program I wrote in 1966 outside the class room for a real-life problem. It was a program that implemented the external functionality of Operation Match, a computer-based dating and match-making service set up in the mid 1960s. I wrote it for my synagogue youth group in order that it could have a dance in which each person's date for the dance was that picked by a variation of the Operation Match software. The dance and the software were called "Operation Shadchan".² I got a hold of the questionnaire for Operation Match, which was being used to match a new client of one gender with previously registered clients of the opposite gender. Each client filled out the form twice, once about him or herself and the second time about his or her ideal mate. For each new client, the data for the client would be entered. Then the software would find all sufficiently good matches with the new clients from among the previously registered clients. Presumably, the matches had to be both good and balanced; that is, the total number of questions for which each answered the way the other wanted had to be greater than a threshold and the difference between the number of matches in each direction had to be smaller than another threshold. I adapted this questionnaire for high school purposes. For example, I changed "Do you believe in sex on the first date?" to "Do you believe in kissing on the first date?".³

I then proceeded to write a program. I remember doing requirements analysis at the same time as I was doing the programming in the typical seat-of-the-pants build-it-and-fix-it-until-it-works method of those days:

- discover some requirements,
- code a little,
- discover more requirements,
- code a little more,
- etc, until the coding was done;
- test the whole thing,
- discover bugs or new requirements,
- code some more, etc.

The first requirements were fairly straightforward to identify. Since this matching was for a dance, unlike with Operation Match, each person would be matched with one and only one person of the opposite gender.⁴ Obviously, we had to make sure that in the input set, the number of boys was equal to the number of girls, so that no one would have the stigma of being unmatched by the software. The next requirements were not so easy to identify. Each boy and each girl should be matched to his or her best match. So, I wrote a loop that cycled through each person and for each, cycled through each other person of the opposite gender to find the best match. But whoa! what is a match? Ah, it cannot be

² "Shadchan" is Yiddish for "Matchmaker". The "ch" in "shadchan" is pronounced as the "X" in "TeX". One person attending the dance thought the name of the dance was "Operation Shotgun".

³ Remember, this was during the mid 1960s!

⁴ It was assumed that each person wanted someone of the opposite gender.

just one way. It must be a mutually balanced match. But double whoa! I have to remove from the list of those that can be cycled through in either loop those that have been matched before. But triple whoa! Suppose the best mutual match for someone is among those not considered because that best mutual match has already been matched to someone else, and that earlier match is not as good. Worse than that, suppose that what is left to match with someone are absolute disasters for the someone. This simple algorithm is not so hot.

In those days and at that age, couples in which the girl was taller than the boy was a disaster, especially if the difference was big. Also, it was not considered so good if the girl of a couple were older than the boy. Therefore, to avoid being painted into a disastrous-match corner, I decided to search in a particular order, from hardest-to-find-non-disastrous matches to easiest. That is, I searched for matches for the tallest girls and shortest boys first and the shortest girls and tallest boys last. Presumably the tallest boys get assigned fairly early to the tallest girls and the shortest girls would get assigned fairly early to the shortest boys. I randomly chose the gender of the first person to be matched and alternated the gender in each iteration of the outer loop. To help avoid disastrous matches, I gave extra weight to the height and age questions in calculating the goodness of any potential match. Each "whoa" above represents a scrapping of previously written code in favor of new code based on the newly discovered requirements. Thus, I was discovering requirement flaws and correcting them during coding as a result of what I learned during coding.

The biggest problem I had was remembering all the requirements. It seemed that each thought brought about the discovery of more requirements, and these were piling up faster than I could modify the code to meet the requirements. I tried to write down requirements as I thought of them, but in the excitement of coding and tracking down the implications of a new requirement, which often included more requirements, I neglected to or forgot to write them all down, only to have to discover them again or to forget them entirely.

Basically, programming felt like skiing down a narrow downhill valley with an avalanche following me down the hill and gaining on me.

Nowadays, we follow more systematic methods. However, the basic feelings have not changed. Since then, I have maintained and enhanced a curve fitting application for chemists⁵. I built a payroll system. I have participated in the writing of a collection of text formatting software. I have watched my graduate students develop tools for requirements engineering. I watched my ex-wife build a simulation system. I have watched my ex-wife and former students and friends involved in startups building large CBSs. To me and, I am sure, the others, programming still feels like skiing with an avalanche following closely behind. I see all the others undergoing similar feelings and being as empathic as I am, I get the same skiing feeling. No matter how much we try to be systematic and to document what we are doing, we forget to write things down, we overlook some things, and the discoveries seem to grow faster than the code.

⁵ This was my second system, and it suffered Brooks's second system syndrome [18], as I tried to build a super-duper, all-inclusive, fancy whiz-bang general curve fitting application with all sorts of fantastic options.

The real problem of software engineering is dealing with ever-changing requirements. It appears that no model, method, artifact, or tool offered to date has succeeded to put a serious dent into this problem. I am not the first to say so. Fred Brooks and Michael Jackson, among others, have said the same for years. Let us examine their arguments.

3 THE SEARCH FOR A SILVER BULLET

Some time ago, Fred Brooks, in saying that there is no software engineering silver bullet⁶ [17], classified software issues into the essence and the accidents. The essence is what the software does and the accidents are the technology by which the software does the essence or by which the software is developed. That is, the requirements are the essence, while the language, tools, and methods used are the accidents. He went on to say, “The hardest single part of building a software system is deciding precisely what to build.... No other part of the work so cripples the resulting system if it is done wrong. No other part is more difficult to rectify later.” This quotation captures the essential difficulty with software that must be addressed by any method that purports to alter fundamentally the way we program, that purports to make programming an order of magnitude easier, that purports to be the silver programming bullet we have all been looking for. Heretofore, no single method has put a dent into this essential problem, although all the discovered methods have combined to improve programming by at least an order of magnitude since 1968, the year the term “software engineering” was invented [59]. Bob Glass reviews the data showing the modest improvements of each of a variety of techniques, methods, and models [33].

Moreover, Brooks, based on his experience, predicted that “There is no single development, in either technology nor management technique, which by itself promises even one order-of-magnitude improvement within a decade in productivity, in reliability, in simplicity.” He made this prediction in 1986 when he first published “No Silver Bullet” in *Information Processing '86*. Since we are now well past a decade after 1986, and no such single technology or management technique has appeared, he has been proven correct. He added a slightly stronger, but still conditional, prediction with which I agree. “*I believe the hard part of building software to be the specification, design, and testing of this conceptual construct, not the labor of representing it and testing the fidelity of the representation. We still make syntax errors, to be sure; but they are fuzz compared to conceptual errors in most systems. If this is true, building software will always be hard. There is inherently no silver bullet.*”⁷ Because we will always be building systems at the new frontiers opened by the latest advances in the accidents, I believe that the antecedent of this conditional statement, that conceptual errors are the hardest kind to find, will always be true. Therefore, *I believe that the conclusion will always be true and that there is inherently no silver bullet.*

⁶ This paper, published in *IEEE Computer in 1987* is a reprint of an earlier, less accessible publication in the 1986 *IFIP* proceedings [16], and is in turn reprinted in the currently more accessible 20th Anniversary Edition of *The Mythical Man-Month* [18].

⁷ I owe Brooks an apology. In a previous version of this paper, I had been guilty of propagating a false quotation incorrectly attributed to him,

Make no bones about it. Software productivity has improved since the publication of “No Silver Bullet” in 1986 because of the accumulative effect of all the accidental advances. Ironically, each such advance makes additional advances more and more difficult, because as each accidental difficulty is solved, what is left is more and more purely of essence.

The obvious question is “Why is there no silver bullet, and why can there not be a silver bullet?” The contention of this paper is that every time a new method that is intended to be a silver bullet is introduced, it does make many parts of the accidents easier. However, as soon as the method needs to deal with the essence or something affecting or affected by the essence, suddenly one part of the method becomes painful, distasteful, and difficult, so much so that this part of the method gets postponed, avoided, and skipped. Consequently, the method ends up being only slightly better than no method at all in dealing with essence-borne difficulties.

But, what *is* so difficult about understanding requirements? I mean, it should be possible to sit down with the customer and users, ask a few questions, understand the answers, and then synthesize a complete requirements specification. However, it never works out that way. Michael Jackson, Paul Clements, David Parnas, Meir Lehman, Bennet Lientz, Burton Swanson, and Laszlo Belady explain why.

4 REQUIREMENTS CHANGE

Michael Jackson, in his Keynote address at the 1994 International Conference on Requirements Engineering [42] said that two things are known about requirements:

1. They will change.
2. They will be misunderstood.

The first implies that a CBS will always have to be modified, to accommodate the changed requirements. Even more strongly, *there ain't no way that requirements are not gonna change*, and there is as much chance of stopping requirements change as there is stopping the growth of a fully functioning and heartily eating teenager. The second implies that a CBS will always have to be modified, to accommodate the changes necessitated by better understanding, as the misunderstandings are discovered. Clements and Parnas describe how difficult it is to understand everything that might be relevant [62].

Meir Lehman [50] classifies a system that solves a problem or implements an application in some real world domain as an E-type system. He points out that once installed, an E-type system becomes inextricably part of the application domain so that it ends

namely the claim that he said, “No major improvement in the software engineering area will ever appear.” Brooks correctly pointed out he had never made such a sweeping assertion. I had seen the quotation somewhere on the Web, and in an effort to avoid potential transcription errors, I had copied and pasted that false quotation into my paper. I did put it into my mind to check against the original document that I got the quotation correct. However, in the confusion of writing a paper, teaching classes, and carrying out administrative duties at the university, I apparently did not do this double check and ended up believing that I had. I am very sorry. In any case, I had no intent to mislead. After all, I proudly sent this paper to Brooks for his comments. Also, I personally agreed with the supposed quote. However, it must be *I* who takes any blame should the sweeping prediction not come true.

up altering its own requirements.

Certainly, not all changes to a CBS are due to requirement changes, but the data show that a large majority of them are. Bennett Lientz and Burton Swanson found that of all maintenance of application software, 20% deal with correcting errors, and 80% deal with changing requirements. Of the requirement changes, 31% are to adapt the software to new platforms, and 69% are for perfecting the software, to improve its performance or to enhance its functionality [53].

Laszlo Belady and Meir Lehman observed the phenomenon of eventual unbounded growth of errors in legacy programs that were continually modified in an attempt to fix errors and enhance functionality [7,8]. That is, as programs undergo continual change their structure decays to the point that it is very hard to add something new or change something already there without affecting seemingly unrelated parts of the program in a way that causes errors. It can be difficult even to find all the places that need to be modified. The programmers make poor guesses and the program, if it even runs, behaves in strange and unpredicted ways.




Figure 1: Belady-Lehman Graph

They modeled the phenomenon mathematically and derived a graph like that of Figure 1, showing the expected number of errors in a program as a function of time, as measured by the ordinal numbers of releases during which modifications are made. In practice, the curve is not as smooth as in the figure, and it is sometimes necessary to get very far into the upswing before being sure where the minimum point is. The minimum point represents the software at its most bug-free release. After this point, during what will be called the *Belady-Lehman (B-L) upswing*, the software's structure has so decayed that it is very difficult to change anything without adding more errors than have been fixed by the change.

The alternative to continuing on the B-L upswing for a CBS is to roll back to the best version, that is, the version that existed at the minimum point. Of course, rolling back assumes that all versions have been saved. All the bugs in this best version are declared to be features, and no changes are ever made in the CBS from then on. Usually not changing a CBS means that the CBS is dead, that no one is demanding changes because no one is using the software any more. However, some old faithful, mature, and reliable programs e.g. *cat* and other basic UNIX applications, vi,

and *ditroff*,⁸ have gone this all-bugs-are-features route. The user community has grown to accept, and even, require that they will never change. If the remaining bugs of the best version are not acceptable features or the lack of certain new features begins to kill usage of the CBS, then a new CBS has to be developed from scratch to meet all old and new requirements, to eliminate bugs, and to restore a good structure to make future modifications possible. Another alternative that works in some special cases is to use the best version as a feature server for what it can do well and to build a new CBS that implements only the new and corrected features and has the feature server do the features of the best version of the old CBS.

The time at which the minimum point comes and the slopes of the curves before and after the minimum point vary from development to development. The more complex the CBS is, the steeper the curve tends to be. Moreover, most of the time, for a carefully developed CBS, the minimum point tends to come in later releases of that CBS. However, occasionally, the minimum point is passed during the development of the first release, as a result of extensive requirements creep during that initial development. The requirements have changed so much since the initial commitment to architecture that the architecture has had to be changed so much that it is brittle. It has become very hard to accommodate new or changed requirements without breaking the CBS. Sometimes, the minimum point is passed during the initial development as a result of code being slapped together into the CBS with no sense of structure at all. The slightest requirements change breaks the CBS.

5 PURPOSE OF METHODS

One view of software development methods is that each method has as its underlying purpose to tame the B-L graph for the CBS developments to which it is applied. That is, each method tries to delay the beginning of the B-L upswing or to lower the slope of that B-L upswing or both. For example, Information Hiding [61] attempts to structure a system into modules such that each implementation change results in changing only the code, and not the interface, of only one module. If such a modularization can be found, then all the code affecting and affected by any implementation change is confined to one module. Thus, it is easier to do the modifications correctly and without adversely affecting other parts of the system. Hence, arrival at the minimum point is delayed and the slope of the upswing is reduced.

In this sense, each method, if followed religiously, works. Each method provides the programmer a way to manage complexity and change so as to delay and moderate the B-L upswing. However, each method has a catch, a fatal flaw, at least one step that is a real pain to do, that people put off. People put off this painful step in their haste to get the software done and shipped out or to do more interesting things, like write more new code. Consequently, the software tends to decay no matter what. The B-L upswing is inevitable.

What is the pain in Information Hiding? Its success in making future changes easy depends on having identified a right

⁸ Well, at least I think so! One great thing about these programs that have not been modified since the 1980s is that their speed doubles every 18 months!

decomposition. If a new requirement comes along that causes changes that bridge several modules, these changes might very well be harder than if the code were more monolithic, simply because it is generally easier on tools to search within one, even big, module than in several, even small, modules [25, 36]. Moreover, future changes, especially those interacting with the new requirement, will likely cross cut [45] module boundaries.⁹ Consequently, it is really necessary to restructure the code into a different set of modules. This restructuring is a major pain, as it means moving code around, writing new code, and possibly throwing out old code for no externally observable change in functionality. It is something that gets put off, causing more modifications to be made on an inappropriate structure.

The major irony is that the reason that the painful modifications are necessary is that the module structure no longer hides all information that should be hidden. Because of changes in requirements, there is some information scattered over several modules and exposed from each of these modules. Note that these changes are requirements changes rather than implementation changes, which continue to be effectively hidden. The painful modifications being avoided are those necessary to restore implementation information hiding, so that future implementation changes would be easier. Without the painful changes, all changes, both implementation and requirements-directed, will be painful. This same pain applies to any method based on Information Hiding, such as Object-Oriented Programming.

In the subsequent sections, each of a number of models, methods, and tools is examined to determine what its painful step is. In some cases, the whole model, method, or tool is the pain; in others, one particular step is the pain. No matter what, when people use these models, methods, and tools, they tend to get stopped by a painful step.

Despite the rhetoric, please understand that I am not opposed to any of the methods I am describing below. I have used some of them successfully. I am merely trying to show where each method becomes painful to use. A mature software engineer would continue to use them even when they are painful.

6 DEVELOPMENT MODELS AND GLOBAL METHODS

This section considers several development models and general programming methods to identify their painful parts. The set chosen is only a sampling. Space limitations and reader boredom preclude covering more. It is hoped that after reading these, the reader is able to identify the inevitable pain in his or her own favorite model or method. In fact, for many models and methods, the pain lies in precisely the same or corresponding activities, all in response to requirements change. The models covered are the Build-and-Fix Model, the Waterfall Model and Requirements Engineering. The general methods covered are Structured Programming, Extreme Programming, Program Generation, Rapid

⁹ For an example that requires minimum explanation, consider two independently developed modules, for implementing unbounded precision *integer* and *real* arithmetic. Each can have its implementation change independently of the other. However, if the requirement of intertype conversion is added, suddenly the two modules have to be programmed together with representations that allow conversion.

Prototyping, and Formal Methods.

6.1 Build-and-Fix Model

The non-method that I applied when I built the Operation Shadchan program was the build and fix model. As indicated by Steve Schach [68], the basic cycle is:

1. Build the first version of the software.
2. Modify the software until client is satisfied.
3. Operate program until a problem is found; when a problem is found, go back to Step 2.

In this case, there is actually no pain, unless the customer's dissatisfaction, the repeated difficulties with the software, or the feeling of skiing in front of an avalanche gives the programmer pain. On the other hand, the model is not purported to be a silver bullet, and the B-L upswing can start almost immediately and can be rather steep, unless the program is small enough to be written entirely by one person.

6.2 Waterfall Model

The waterfall model [67] is an attempt to put discipline into the software development process by forcing understanding and documentation of the requirements before going on to design, by forcing understanding and documentation of design before going on to coding, by forcing thorough testing of the code while coding each module, etc. The model would work if the programmers could understand a stage thoroughly and document it fully before going on to the next stage [62]. However, understanding is difficult and elusive, and in particular, documentation is a pain. The typical programmer would prefer to get on to coding before documenting. Consequently, some programmers consider the whole model to be a pain, because it tries to force a disciplined way of working that obstructs programmers from getting on with the coding in favor of providing seemingly endless, useless documentation. However, even for the programmers who do believe in discipline, the waterfall becomes a pain in any circumstance in which the waterfall cannot be followed, e.g., when the full requirements are learned only after significant implementation has been carried out or when there is significant repeated backtracking, when new requirements are continually discovered.

6.3 Structured Programming

I can recall the first systematic programming method I learned, used, and taught, from 1973 until it fell out of fashion in the mid 1980s, namely, Structured Programming or Stepwise Refinement [75, 23]. After having done build-it-and-fix-it programming for years, Structured Programming appeared to be the silver bullet, at last, a way to put some order into the chaotic jumble of thoughts that characterized my programming, at last, a way to get way ahead of or to the side of the avalanche that was coming after me all the time.

In fact, I found that Structured Programming did work as promised for the development of the first version of any CBS. If I used the clues provided by the nouns that appeared in more than one high-level statement [9, 11], Structured Programming did help me keep track of the effects of one part of the program on another. It

did help me divide my concerns and conquer them one-by-one, without fear of forgetting a decision I had made, because these decisions were encoded in the lengthy, well-chosen names of the abstract, high-level statements that were yet to be refined. Best of all, the structured development itself was an ideal documentation of the structure of the program.

However, God help me if I needed to change something in a program developed by stepwise refinement, particularly if the change was due to an overlooked or changed requirement. I was faced with two choices:

1. Patch the change into the code in the traditional way after a careful analysis of ripple effects; observe that the documented structured development assists in finding the portions of the code affected by and affecting the change.
2. Redo the entire structured development from the top downward, taking into account the new requirement and the old requirements, all in the right places in the refinement.

The problems with the first choice are that:

1. no matter how careful I was, always some ripple effects were overlooked, and
2. patching destroys the relation between the structured development and the code, so that the former is no longer accurate documentation of the abstractions leading to the code. This discrepancy grows with each change, thus hastening the onset of the B-L upswing.¹⁰ Thus, the new code contradicts the nice structure imposed by the structured development. Moreover, the names of the high level abstractions that get refined into code no longer imply their refinements.

Therefore, the correct alternative was the second, to start from scratch on a new structured development that recognizes the modified full set of requirements. However, then the likelihood is that the new code does not match the old code. Most structured programmers do this redevelopment with an eye to reusing as much as possible. That is, whenever one encounters a high-level statement with identical name and semantics as before, what it derived before is put there. However, the programmer must be careful to use the same high-level statements as refinements whenever possible, and he or she must be careful that in fact the same semantics is wanted as before and that he or she is not self deluding to the point of falsely believing a high-level statement has the same semantics as before.

This second alternative turned out to be so painful that I generally opted to the first alternative, patching the code and thus abandoned the protection, clarity, and documentation offered by Structured Programming.

About that time, I learned the value of faking it. Do the modification by patching up the code, and then go back and modify the original structured development to make it look like the new program was derived by Structured Programming. However, faking it was also painful, and soon, unless I was required to do so for external reasons, e.g., preparing a paper for publication or a contract, I quickly stopped even faking it. Adding to the pain of faking it was the certainty of not doing the faking perfectly and being caught. It was only later that David Parnas and Paul Clements legitimized faking it [62] and relieved my guilt.

¹⁰ It is clear that there is a discrepancy, because if the abstractions had derived the new code, then the new code would have been there before.

6.4 Requirements Engineering

The basic premise of requirements engineering is spend sufficient time up front, before designing and coding, to anticipate all possible requirements and contingencies, so that design and coding consider all requirements and contingencies from the beginning [13, 60]. Consequently, fewer changes are required, and the B-L upswing is both delayed and moderated. Data show that errors found during requirements analysis cost one order of magnitude less to fix than errors found during coding and two orders of magnitude less to fix than errors found during operation [15]. These economics stem from the very factors that cause the B-L upswing, namely, the fact that in a full running program, there is a lot more code affecting and affected by the changes necessary to fix an error than in its requirements specification. There are data and experiences that show that the more time spent in requirements engineering, the smoother the implementation steps are [31, 24], not only in software engineering, but also in house building [14, 73]. When the implementation goes smoother, it takes less time, it is more predictable, and there are fewer bugs.

However, for reasons that are not entirely clear to me,¹¹ a confirmed requirements engineer, people seem to find haggling over requirements a royal pain. They would much rather move on to the coding, and they feel restless, bored, or even guilty when forced to spend more time on the requirements. A boss with his or her eyes focused on an unrealistically short deadline does not help in this respect. I would bet that Arnis Daugulis [24], his colleagues, and bosses at Latvenergo felt the pain of not being able to move on to implementation, even though in the end, they were happy that they did not move on to implement the first two requirements specifications, which were, in retrospect, wrong.

The pain is exacerbated, and is felt even by those willing to haggle the requirements, because the requirements engineer must make people discover requirements by clairvoyance rather than by prototyping. The pain is increased even more as the backtracking of the waterfall model sets in, as new requirements continue to be discovered, even after it was thought that all requirements had been found.

There appear to be at least two ways of carrying out RE in advance of CBS design and implementation, what are called in the agile software development community [1] “Big Modeling Up Front (BMUF)” [3] and “Initial Requirements Up Front (IRUF)” [2]. They differ in the ways they treat the inevitable requirements changes.

In BMUF, the CBS developers try to create comprehensive requirement models for the whole system up front, and they try to specify the CBS completely before beginning its implementation. They try to get these models and specifications fully reviewed, validated, agreed to, and signed off by the customer and users. Basically, the developers try to get the requirements so well understood that they can be frozen, no matter how painful it is. However, as demonstrated in Section 4, the requirements continue

¹¹ Then again, I always read the manual for an appliance or piece of hardware or software completely before using the appliance or piece. I return the appliance or piece for a refund if there is *any* disagreement between what the manual says and the appliance or piece does, even if it is in only format of the screen during the set up procedure.

to change as more and more is learned during implementation. To stem the tide of requirements change, the powers with vested interest in the freezing of the requirements create disincentives for changes, ranging from contractual penalties against the customer who demands changes to heavy bureaucracy, in the form of a change review board (CRB). For each proposed change, the CRB investigates the change's impact, economic and structural. For each proposed change, the CRB decides whether to reject or accept the change, and if it accepts the change, what penalties to exact. Of course, the CRB requires full documentation of all requirements models and specifications and up-to-date traceability among all these models and specifications.¹²

Agile developers, on the other hand, expect to be gathering requirements throughout the entire CBS development. Thus, they say that we should “embrace change” [2]. We should explore the effects of a proposed change against the organizations business goals. If the change is worth it, do it, carrying out the required modifications, including restructuring. If the change isn't worth it, don't do it [40]. It's that simple.

The objective of agile processes is to carry out a CBS “development project that

- (1) focuses and delivers the essential system only, since anything more is extra cost and maintenance,
- (2) takes into account that the content of the essential system may change during the course of the project because of changes in business conditions,
- (3) allows the customer to frequently view working functionality, recommend changes, and have changes incorporated into the system as it is built, and
- (4) delivers business value at the price and cost defined as appropriate by the customer.” [69]

Accordingly, agile developers get the IRUF, with all the stakeholders, i.e., customers, users, and developers, participating actively at all times, in which as many requirements as possible are gathered up front from all and only stakeholders. The goals of getting the IRUF are [2]

1. to identify the scope of the CBS being built,
2. to define high-level requirements for the CBS, and
3. to build consensus among stakeholders as to what the requirements imply.

The IRUF session ideally is as short as a few hours, but it can stretch into days and even weeks in less than ideal circumstances, such as not all stakeholders being in one place, or particularly tricky requirements. Then come several days of modeling sessions to produce a full set of models of the CBS as conceived by the IRUF. These models include use cases, which are eminently suitable for discussions between users and developers.

Following this modeling, the requirements are ranked by priority by all stakeholders. Business goals are taken into account during this ranking process, which typically requires about a day. Detailed modeling of any requirement takes place only during the

¹² The agile development community says that traceability is a waste of resources. The cost of keeping trace data up to date must be balanced against the cost of calculating the trace data when they are needed to track down the ripple effects of a proposed change. The community believes that updating is so much more frequent than tracing that the total resources spent in continually updating outstrips the resources spent in occasional tracing.

beginning of the iteration during which it is decided to implement that requirement.

Notice that BMUF and IRUF differ in their approaches to dealing with the relentless, inevitable requirements changes. BMUF tries to anticipate all of them. IRUF does not; it just lets them come. BMUF is considered as not having succeeded totally if it fails to find a major requirement. The pain of dealing with the change with BMUF is felt during the RE process. IRUF practitioners embrace the change and decide on the basis of business value whether or not to do the change. The pain of dealing with the change with IRUF is felt in the re-implementation necessitated by the change, e.g., in the refactoring that can be necessary. See Section 6.5 for details about refactoring pain. Ultimately, neither approach can *prevent* a new requirement from appearing.

6.5 Extreme Programming

Extreme Programming (XP) [6] argues that the preplanning that is the cornerstone of each the various disciplined programming methods, such as the Waterfall model, documentation, requirements engineering, is a waste of time. This preplanning is a waste of time, because most likely, its results will be thrown out as new requirements are discovered. XP consists in simply building to the requirements that are understood at the time that programming commences. However, the requirements are given, not with a specification but with executable test cases. Unfortunately, these test cases are not always written because of the pain of writing test cases in general and in particular before it known what the code is supposed to do [56]. (See Section 7.2 for more on testing.) During this programming, a number of proven, minimum pain, and lightweight methods are applied to insure that the code that is produced meets the requirements and does so reliably. The methods include continual inspection, continual testing, and pair programming. Thus, the simplest architecture that is sufficient to deal with all the known requirements is used without too much consideration of possible changes in the future and making the architecture flexible enough to handle these changes. It is felt that too much consideration of the future is a waste, because the future requirements for which it plans for may never materialize and an architecture based on these future requirements may be wrong for the requirements that do come up in the future.

What happens when a requirement comes along that does not fit in the existing architecture? XP says that the software should be refactored. Refactoring consists in stepping back, considering the new current full set of requirements and finding a new simplest architecture that fits the entire set. The code that has been written should be restructured to fit the new architecture, as if it were written with the new architecture from scratch. Doing so may require throwing code out and writing new code to do things that were already implemented. XP's rules say that refactoring should be done often. Because the code's structure is continually restructured to match its current needs, one avoids having to insert code that does not match the architecture. One avoids having to search widely for the code that affects and is affected by the changes. Thus, at all times, one is using a well-structured modularization that hides information well and that is suited to be modified without damaging the architecture. Thus, the B-L upswing is

delayed and moderated.

However, refactoring, itself, is painful [28]. It means stopping the development process long enough to consider the full set of requirements and to design a new architecture. Furthermore, it may mean *throwing out perfectly good code whose only fault is that it no longer matches the architecture*, something that is very painful to the authors¹³ of the code that is changed. Consequently, in the rush to get the next release out on time or early, refactoring is postponed and postponed, frequently to the point that it gets harder and harder. Also, the programmers realize that the new architecture obtained in any refactoring will prove to be wrong for some future new requirements. Ironically, the rationale for doing XP and not another method, is used as an excuse for not following a key step of extreme programming.

Indeed, Elssamadisy and Schalliol report on an attempt to apply a modified XP approach to a software development project that was larger than any previous project to which they had applied XP, with great success. They use the term “bad smells” to describe symptoms of poor practice that can derail XP from its usual swift track [28]. Interestingly, each of these bad smells has the look, feel, and odor of an inevitable pain.

1. In XP, to ensure that all customer requirements are met, a customer representative has to be present or at least available at all times and he or she must participate in test case construction. The test cases, of course, are written to test that the software meets the requirements. Elssamadisy and Schalliol report that over time, the customer representatives began to “refrain from that ‘toil’” and began to rely on the test cases written by the developers. Consequently, the test cases ceased to be an accurate indication of the customer’s requirements, and in the end, the product failed to meet the customer’s requirements even though it was running the test cases perfectly. It appears to me that participation in the frequent test case construction was a pain to the customer that was caused by the relentless march of new requirements.
2. XP insists on the writing of test cases first and on delivery of running programs at the ends of frequent development iterations in an attempt to avoid the well-known horror of hard-and-slow-to-fix bug-ridden code that greets many software development teams at delivery time. However, as a deadline for an iteration approaches, and it is apparent that programming speed is not where it should be, managers begin to excuse developers from having to write test cases; developers refrain from just barely necessary refactorings; testers cut back on the thoroughness of test cases; and developers fail to follow GUI standards and to write documentation, all in the name of sticking to the iteration schedule. Speed is maintained, but at the cost of buggy code. It appears to me that all curtailed activities were painful in XP, just as they are in other methods.
3. Elssamadisy and Schalliol report that there was a routine of doing a particular function, called **unbooking** and **rebooking** of a **lease** (URL). The first version of the routine did URL in the only way that was required by the first story. As a new story required a different specific case of URL, the required code was patched in to the routine as a special case. With each

such new specific case of URL, the code became patchier and patchier; it became more and more apparent that a more and more painful refactoring was needed so that each specific case of URL is a specialization of a more generic, abstract URL. Elssamadisy and Schalliol

were the ones who got stuck with fixing the problem. Also, because one of them is a big whiner, he kept complaining that “we knew this all along—something really stank for iterations and now I’m stuck with fixing it.” The sad truth is that he was right. Every developer after the initial implementation knew that the code needed to be refactored, but for one reason or another ..., they never made the refactoring.

This type of smell has no easy solution. The large refactoring had to be made because the inertia of the bad design was getting too high (footnote: Look ahead design would have been useful here also.) By taking the easy road from the beginning and completely ignoring the signs, we coded ourselves into a corner. So the moral of the story is this: when you find yourself making a large refactoring, stop and realize that it is probably because that [sic] you have skipped many smaller refactorings [28].

Is this ever a description of pain? It comes in a method which has been carefully designed to avoid many painful activities that programmers dislike.

The claim by agile developers is that the requirements to be addressed in the first iterations are those with the biggest business values. Any other requirements to be considered later are less important, and are even less likely to survive as originally conceived or even at all, due to the relentless change of requirements. Thus, refactoring becomes less and less important with each iteration.

However, this claim presupposes that the initial set of requirements, the IRUF, is so comprehensive that no new important requirements, forcing a major, painful refactoring, will ever be discovered. It’s hard for me to imagine, and it runs counter to my experience that, any requirements modeling that is not the result of BMUF be so successful so as to effectively make refactoring completely unnecessary. Certainly, even if BMUF is done, it’s hard not to ever need refactoring. Thus in the end, the potential for pain is there.

6.6 Program Generation and Domain-Specific Approaches

One approach to simplifying programming is to use a program generator to write a complete program from a declarative specification of the program. For example, one can use a compiler-compiler to write most of a compiler from some syntax and semantic specifications for the language being compiled. These program generators do significantly simplify the process of writing a complete program. However, program generators are generally for application domains that are well enough understood that major requirement surprises that force refactoring of the

¹³ Remember that it’s *pair* programming.

system's architecture, are unlikely. In more understood domains, programs are really more manufactured than they are programmed. In less understood domains, domain-specific approaches provide enough structure, a stable architecture that the domain has become a true engineering discipline, like aircraft, automobile, and bridge engineering. In these domains, each new instance is a perturbation of all previous instance. There are surprises during construction, and significant ones too, but these surprises are still rare enough to force a wholesale change the architecture of the system.

However, even program generators are not without pain. I recall that my third major software development project, for a summer job, was to develop a set of report programs operating on a single employee database. All the programs in the set were to be generated by RPG, the old Report Program Generator. To produce any program in the set, I specified the input data it needed in the database and the format and formulae for the output data, called the report. RPG would read this specification and write a program in assembly language which then had to be assembled. I was glad that I did not have to modify the generated program, which looked nothing like what a human being would write. To change a program, all I had to do was change its specification and to submit the new specification to RPG, which generated the new program. What could be simpler?

Nevertheless, after a while, after the number of related programs grew to a number, n , about 10, each change of data format, each new field, each deleted field, etc. became a nightmare as the ripple effects reverberated to all n programs. I felt like I was skiing ahead of an avalanche even though it was a summer job in an area noted for its 90s (90°F and 90% humidity) weather.¹³ If I forgot to make a change or made the wrong change in one or more of the specifications, the next payroll would be a mess.

Basically, a program generator, and for that matter all automatic programming schemes [5], just move programming back to earlier in the lifecycle to a specification that is at a higher level than the normal program. However, even a high-level specification, because it must be internally consistent, gives pain when it must be modified.

6.7 Rapid Prototyping

Rapid prototyping [4,21] is offered as a means to identify requirements, to try out various ideas, to do usability testing, etc. To the extent that it succeeds, it delays and moderates the B-L upswing. The pain is to throw the prototype out and to start anew when implementation starts. Often, in the face of an impending implementation deadline, instead, the prototype is used as a basis for the first implementation, thus ratifying poor design decisions that were made for expedience in putting together the prototype rapidly. The need to throw out the prototype and start all over with a systematic development is at least the need to refactor if one is following XP. When the prototype's architecture is used for the production software, the very first production version is brittle, hard to modify, and prone to breaking whenever a change is made.

¹³ On the other hand, in those days, computer rooms were so fiercely airconditioned that I had to wear my skiing sweater in the machine room.

6.8 Formal Methods

There are a number of formal program development methods, each of which starts with a formal specification of the CBS to be built [74,22,34,37,44]. Once a formal specification has been written it can be subjected to verification that it meets higher-level formally stated requirements, such as security criteria [43]. Any code written for the implementation can be verified as correct with respect to the specification [41,27]. Sometimes implementing code can be generated directly from the specification [5].

Some consider the mere writing of a formal specification a pain. Some consider writing such a specification to be fun, but consider all the verification that needs to be done to be a pain. Some consider this verification to be fun. However, my experience is that everyone considers the work dealing with changed requirements to be a pain. The specification has to be changed with the same difficulties as changing code, that is, of finding all other parts of the specification affecting or affected by the change at hand.

The worst of all is the necessary reverification. Since requirements are inherently global, formal statements about them and proofs about the formal statements are inherently global as well. Potentially every proof needs to be redone, because even though a theorem clearly still holds, the old proof may use changed formulae. Some have tried to build tools to automatically redo proofs in the face of changing specifications [57,12]. The pain is enough to drive one to adopt a lightweight method in which only the specification is written and compiler-level checks are done and no verification is done [10].

7 SPECIFIC METHODS AND TECHNIQUES

Besides models and methods for the overall process of software engineering, there are a large variety of specific methods and techniques each of which is for attacking one specific problem of software engineering. Even these small methods have their pains, exactly where they brush up against changes. Again, only a small sampling of the available methods are covered in the hope that they are a convincing, representative sample.

7.1 Inspection

The effectiveness of inspection [29] as a technique for finding errors has been documented widely [32]. When applied to early documents, such as requirements or design documents, inspection helps delay and moderate the B-L upswing. However, inspection is a double pain. First, the documents to be inspected must be produced, and we know that documentation itself is a pain. Second, inspection is one of these unpopular activities that are the first to be scrubbed when the deadlines are looming [32]. Roger Pressman quotes Machiavelli in dealing with the unpopularity of inspection.

“... some maladies, as doctors say, at the beginning are easy to cure but difficult to recognize ... but in the course of time ... become easy to recognize but difficult to cure.” Indeed! But we just don't listen when it comes to software. Every shred of evidence indicates that formal technical reviews (for example, inspections) result

in fewer production bugs, lower maintenance costs, and higher software success rates. Yet we're unwilling to plan the effort required to recognize bugs in their early stage, even though bugs found in the field cost as much as 200 times more to correct [64].

7.2 Regression Testing

Regression testing is subjecting a new version of a CBS, which was constructed in response to a detected fault or a requirements change, to *all* previously run tests, in an effort to regress to the faults detected previously by them.¹⁵ Of course, if a requirement has changed, then the expected response to some of the old test input may be different from before. Moreover, new tests testing the new features have to be added to the set for now and for future regression tests.

The necessity of running all the old tests arises from the danger that even the tiniest change in one place of the software have an unpredicted effect on another part of the software. These thorough tests of the whole system are necessary to try to expose these ripple effects.

The whole process is such a pain. Even if the running of the old tests and the comparison of the actual and expected results are automated, examination of reported differences is no fun. Differences mean having to go back and do more work. They show inadequacies in one's programming. *The regression tests take time that could be better spent on something useful like more coding! Besides which, it's obvious that my teeny, little ol' one-line change did not affect anything else.* Thus, there is a tendency to skip this step after each change and to postpone it until later or until never. The consequence of the postponement is that when the tests are run and an error is exposed, it is not clear which of the many little changes since the last regression tests caused the error. Had a full regression been done after each change, the source of any error found would be much clearer; it has to be at least partially due to the last change made.

7.3 The Daily Build

In a situation in which lots of programmers are continually modifying relatively but not completely independent code modules in a single CBS to fix bugs, improve performance, add new features, etc., a common practice is the process called the *daily build* [55]. It works best when there is tool support for version control, including commands for checking a module out, checking a module in, and identifying potential conflicts between what is checked in and what is already in.

Each programmer gets a problem, perhaps a bug to fix, performance to improve, or a feature to add, and he or she plans his or her attack. When ready, he or she downloads the latest versions of each module that needs to be changed. He or she makes the changes, he or she tests and regression tests the changed modules, he or she tests and regression tests the whole changed system, and when he or she is satisfied that the changes will work, he or she checks the modules back in.

¹⁵ Remember that the purpose of testing is to *find* errors, not to show that there are none [58].

The fun comes when each day, whatever complete collection of modules is there is compiled into a running version of the system and is tested and regression tested. If the system passes all the test, everything is fine. If not, whoever made the last change has to fix the problems, perhaps in consultation with other programmers, particularly of the modules affected by his or her changes. Consequently, the incentive is for each programmer to work quickly and to verify that the version of each module that is there just before a check in is what was checked out. If not, he or she should certainly feel compelled to re-implement and test his or her changes on that latest version. If the changes made by others are truly independent, this reimplementation is straightforward, and consists of folding in the independent changes identified with the help of *diff*. Obviously, the faster he or she works after checking out the modules to be changed, the less likely he or she will find them changed by others at checkin. Note that two programmers can work at cross purposes. However, if they notice that they are working on the same modules, they are encouraged to work together to ensure that their changes do not conflict.

The testing, regression testing, the checking in, and possible rework are all real pains. Most people would just as soon dispense with them to get on to the real work, programming. However, in this case, the social and work costs of failing to do these activities is very high. Basically, the one whose checkin caused the build not to work correctly is the turkey of the day to be hung by his or her thumbs, and he or she has to spend considerable time to fix a version that has multiple programmers' changes.

7.4 Open Sourcing

Open sourcing [65] is a world-wide application of the daily build, compounded from daily to continually. The main advantage is that the software has a world-wide legion of merciless inspectors and bug fixers. Whoever sees the source of a problem is encouraged by the promise of fame to fix it. "Given enough eyeballs, all bugs are shallow." The pain for each programmer is integrating his or her update in the face of almost continual change. The biggest pain is for the manager of the whole project, e.g., Linus Torvalds for Linux. Can you imagine what he has to go through to reconcile different fixes of a problem, to pick the one that he thinks fits best in a moving target. Compound this problem with the fact that he must deal with several problems at once. Goodness of fit is measured not only by how well it solves the problem but also by how consistent it is with all other changes that are being done. If the manager is not careful, the source could slide into B-L upswing oblivion very quickly.

8 DOCUMENTATION

Almost all methods, including the general waterfall model, require programmers to document their decisions so that all decisions are visible to all persons working on the same CBS development, and even to the persons who make the decisions and later forget them. For example, Information Hiding requires that the interfaces and behaviors of the procedures and containing modules be documented so that users know what the modules are supposed to do without having to see the hidden code. When the documented information is available, the methods work as

intended and the B-L upswing is both delayed and moderated. The kinds of documentation included are:

1. in the waterfall model, the requirements specification, the project plan, the design, the test plan, the comments, etc.,
2. in Structured Programming, the structured development itself,
3. in Information Hiding, besides the interfaces and behaviors, for each module, the secret that it hides; this secret is the encapsulated design information of the module [39],
4. in traceability schemes, the links between related units of the artifacts, which can be any other kind of documentation and the code itself [35], and
5. in Requirements Engineering, the requirements specification, the lexicon giving the problem vocabulary [52], the scenarios describing how the CBS is supposed to be used [71, 20], and sometimes a network of issues, positions, and arguments for capturing rationale [19].

A number of authors, including S. D. Hester, D. L. Parnas, and D. F. Utter [39], have gone so far as to advocate using the documentation that one should write anyway as a design medium.

However, documentation itself is painful. It requires that people stop the analysis, design, or coding work that they are doing and to write their ideas down on the spot. If they do not stop to document, they forget the details, and if and when there is time later to document, a lot of the information that should have gone into the documentation has been forgotten. Often, *there is no time to document*. The shipping deadline is looming. So, the documentation never happens. If any of it happens, it quickly gets out of date as the program continues to change, and people continue to postpone recording decisions. Later, the programmers who must make changes have no information or poor, out-of-date information about the requirements of, the architecture of, the design of, the rationale behind, the test plans for, and the commentary about the code. Since they do not know what documentation to trust, they trust none of it [26]. Therefore, they have a hard time finding the code and tests affecting and affected by the changes. The B-L upswing sets in earlier and is steeper.

Again, the irony is that the very process that the documentation helps is the process that undermines the validity of the documentation.

9 TOOLS AND ENVIRONMENTS

There are a number of environments of tools that help manage CBS development by keeping all artifacts on-line in an editable form, by keeping relationships between the artifacts, by performing a variety of consistency checks and analyses, and by doing other mundane tasks such as running regression tests. Among the information kept on line in such environments are

- all the documentation mentioned in Section 8 and
- configuration information including module dependencies [30] and version control information [70, 66]

Section 8 describes the use and the pain of the documentation. The configuration information allows groups of programmers working on different versions of a program to keep some sanity while being able to work somewhat independently on different modules and still be able to generate the different versions at any time. When two programmers have been working on the same

modules or modules that depend on each other, the system suggests that they meet to discuss the consistency of their changes.

The general pain of all these tools and environments is remembering to use them and keep their information up to date in the pressure of an impending deadline breathing down the developers' necks. For example, the configuration management systems depend on programmers' checking modules out for working on them and on programmers' checking modules back in when the programmers are finished with working on them. They depend on programmers' meeting when the system suggests that they have to meet. When a programmer fails to check a module out, check a module back in, or to meet with another, the system loses control and the completeness and consistency it promises cannot be guaranteed any more. These requirements on the programmers are really not much of a pain, but programmers do occasionally forget to do one step, again in the rush to beat the deadline.

Keeping on-line artifacts up to date is both easier and harder than keeping paper copies up to date. It is easier because it is on line and there is an editor to help do it. It is harder because the tool often requires more precision in keeping them up to date than a human being would demand. Keeping on-line artifacts in an integrated environment is both easier and harder than keeping independent on-line copies up to date. It is easier because in the environment there are often tools that help find effects on other artifacts. It is harder because usually there are many more artifacts in an integrated environment than there would be if artifacts were maintained independently or on paper. Thus there are many more connections between the artifacts that have to be maintained along with the contents of the artifacts.

10 CONCLUSIONS

Thus, it appears that there is no software engineering silver bullet.¹⁶ All software engineering bullets, even those that contain some silver, are made mostly of lead. It is too hard to purify the painful lead out of the real-life software engineering bullet to leave a pure painless silver software engineering bullet.

The situation with software engineering methods is not unlike that stubborn chest of drawers in the old slapstick movies; a shlimazel¹⁷ pushes in one drawer and out pops another one, usually right smack dab on the poor shlimazel's knees or shins. If you find a new method that eliminates an old method's pain, the new method will be found to have its own source of pain.

There cannot be any significant change in programming until someone figures out how to deal, with a lot less pain, with the relentless change of requirements and all of its ripple effects. Perhaps, we have to accept that CBS development is an art and that no amount of systematization will make it less so.

Software engineering is an art, no less than painting. Indeed, the first part of the titles of Don Knuth's books in his series on algorithms for computer programming is *The Art of Computer*

¹⁶ Some have called the feeling that there is no silver bullet pessimistic. However, I view the feeling as optimistic, because it says that software engineering can *never* be automated, that it will always require thinking, creative, human beings. Therefore, we programmers are always assured of jobs!

¹⁷ "Shlimazel" is Yiddish for "one who always suffers bad luck".

Programming [46,47,48]. The fact that software engineering is an art does not minimize its technical roots. A programmer must know the languages, the hardware platforms, the domain areas of his or her software, etc. However, what he or she does with them is very much a matter of art.

Even traditional arts, e.g., painting, have technical aspects. No matter how talented a painter is, if he or she does not know how to work with paints, his or her paintings will not be good. Furthermore, as has been demonstrated by mediocre artists, knowledge of the techniques does not insure good quality art.

Actually software engineering is an art just like mathematics. Both creativity and knowledge are required. Would mathematicians get upset if they were told that it was impossible to turn mathematics into a systematic engineering discipline? Would mathematicians even try?

If we know a domain well enough that architecture-changing requirement surprises are significantly tamed, as for compiler production these days, we can go the engineering route for that domain, to make building software for it as systematic as building a bridge or a building. However, for any new problem, where the excitement of major innovation is, there is no hope of avoiding relentless change as we learn about the domain, the need for artistry, and the pain.

The key concept in the Capability Maturity Model (CMM) [63] is *maturity*. Getting the capability to do the recommended practices is not a real problem. The real problem is getting the maturity to stick to the practices despite the real pain.

As some, including Alexander Egyed and Dino Mandrioli in private conversation, have observed, the psychology of computer programming [72] is coming to play here. Psychologically, we programmers tend to think coding as the real work and we feel disturbed by, pained by, and as wasting our time in, doing anything else, no matter how necessary the anything else is to producing quality software.

As Les Belandy and Jack Goldberg have observed in private communication, designers and builders of large, complex systems in other engineering disciplines, e.g., of aircraft, bridges, buildings, nuclear power plants, etc., do not complain that there is no magic silver bullet for their discipline. What is the difference? Well, the engineers in the other disciplines know that there is no silver bullet. So why do we software engineers search for a silver bullet?

The search for a software engineering silver bullet has been motivated by a belief that there *has* to be or at least there *ought* to be a software engineering silver bullet. After all, why bother looking for something that does not exist? The search is based on the assumption that programming can be made easy or at least systematic.

However, my own experience has been that programming is extremely difficult, of the order of difficulty of theorem proving,¹⁸ and requires much creativity as well as good ol' fashioned sweat of the brow! While certain aspects of programming, namely the production of code snippets that meet very well-understood requirements, i.e., the body of a procedure or loop not more than a

¹⁸ I am reminded of a paper by Manna and Waldinger showing that programming is equivalent to a constructive proof of the existence of code that is partially correct with respect to input-output assertions [54].

few pages long, coming to grips with any other aspect of code, i.e., requirements, architecture, changes to requirements or architecture, that is, the essence, requires much creativity, blood, sweat, and tears. I am certainly not the only one to say so. Also Fred Brooks [17], Richard Gabriel [38], Don Knuth [49], and John Knight, in private communication, have said so. In particular, Knuth said,

What were the lessons I learned from so many years of intensive work on the practical problem of setting type by computer? One of the most important lessons, perhaps, is the fact that SOFTWARE IS HARD.... From now on I shall have significantly greater respect for every successful software tool that I encounter. During the past decade I was surprised to learn that the writing of programs for T_EX and for METAFONT proved to be much more difficult than all the other things I had done (like proving theorems or writing books). The creation of good software demands a significantly higher standard of accuracy than those other things do, and it requires a longer attention span than other intellectual tasks.

One referee suggests that the focus of this paper is on the interface between the essence and the accidents of software. Heretofore, software engineering's focus has been on one or the other. It is necessary to consider also the interface between them and of the feedback loop caused by that interface [51]. It is my belief that the interface between the essence and the accidents is of essential difficulty.

However, a more perplexing question is how we, in computer science and software engineering ever go into this search for a silver bullet business, which continues to this very day, as indicated by the continuing claims of how method X will improve your programming productivity by 100% or some such large number? How did we ever get into this believe that programming is easy or at least systematizable? Perhaps our fields' origins in Mathematics is responsible for this delusion in a perversely negative way. Some mathematicians, or least some that I have known, tend to look down at programming as not being as intellectually substantial as doing mathematics, as proving theorems, which is regarded as the ultimate in creativity. This belief has been the basis for denials of tenure to people who do only software, no matter how substantial, innovative, and creative it may be. This belief has been the basis for disowning of nascent software engineering programs. Many of us in software engineering have had to fight this shortsightedness. The irony is that while fighting this belief in others, we possibly have adopted it, as the basis for our belief that programming ought to be easy or at least systematizable.

Perhaps now it is clear why I no longer get excited over any new language, development model, method, tool, or environment that is supposed to improve programming.¹⁹ It is clear also why I think that the most important work is that addressing requirements, changes, and the psychology and sociology of programming.

¹⁹ One reviewer offered an alternative explanation, saying "It's because you are getting old!"

ACKNOWLEDGMENTS

I am grateful to Matt Armstrong, Les Belady, Barry Boehm, Fred Brooks, Alexander Egyed, Bob Glass, Jack Goldberg, John Knight, Andrew Malton, Dino Mandrioli, Daniel Morales Germán, John Mylopoulos, Torsten Nelson, Peter Neumann, David Parnas, James Robertson, Steve Schach, Davor Svetinović, and the anonymous referees of previous versions of this paper for useful discussions about specific methods and other issues and careful and thoughtful comments about the paper, both favorable and not so favorable. I was supported in part by NSERC grant NSERC-RGPIN227055-00.

REFERENCES

- [1] “Principles: The Agile Alliance”, The Agile Alliance (2001), <http://www.agilealliance.org/>.
- [2] Ambler, S.W., “Agile Requirements Modeling”, The Official Agile Modeling (AM) Site (2001), <http://www.agilemodeling.com/essays/agileRequirementsModeling.htm>.
- [3] Ambler, S.W., “Agile Software Development”, The Official Agile Modeling (AM) Site (2001), <http://www.agilemodeling.com/essays/agileSoftwareDevelopment.htm>.
- [4] Andriole, S.J., “Fast Cheap Requirements: Prototype or Else!”, *IEEE Software* **14**(2), pp. 85–87 (March 1994).
- [5] Balzer, R.M., “Transformational Implementation: An Example”, *IEEE Transactions on Software Engineering* **SE-7**(1), pp. 3–14 (January 1981).
- [6] Beck, K., *Extreme Programming Explained: Embrace Change*, Addison-Wesley, Reading, MA (1999).
- [7] Belady, L.A. and Lehman, M.M., “A Model of Large Program Development”, *IBM Systems Journal* **15**(3), pp. 225–252 (1976).
- [8] Belady, L.A. and Lehman, M.M., “Program System Dynamics or the Metadynamics of Systems in Maintenance and Growth”, in *Program Evolution*, ed. M.M. Lehman and L.A. Belady, Academic Press, London, UK (1985).
- [9] Berry, D.M., “An Example of Structured Programming”, *UCLA Computer Science Department Quarterly* **2**(3) (July 1974).
- [10] Berry, D.M., “The Application of the Formal Development Methodology to Data Base Design and Integrity Verification”, *UCLA Computer Science Department Quarterly* **9**(4) (Fall 1981).
- [11] Berry, D.M., “Program Proofs Produced Practically”, *Tutorial Notes of Fifth International Conference on Software Engineering* (March 1981).
- [12] Berry, D.M., “An Ina Jo Proof Manager for the Formal Development Method”, in *Proceedings of Verkshop III, Software Engineering Notes* (August 1985).
- [13] Berry, D.M. and Lawrence, B., “Requirements Engineering”, *IEEE Software* **15**(2), pp. 26–29 (March 1998).
- [14] Berry, D.M., “Software and House Requirements Engineering: Lessons Learned in Combating Requirements Creep”, *Requirements Engineering Journal* **3**(3&4), pp. 242–244 (1998).
- [15] Boehm, B.W., *Software Engineering Economics*, Prentice-Hall, Englewood Cliffs, NJ (1981).
- [16] Brooks, F.P. Jr., “No Silver Bullet”, in *Information Processing '86*, ed. H.-J. Kugler, Elsevier Science Publishers B.V. (North-Holland), IFIP (1986).
- [17] Brooks, F.P. Jr., “No Silver Bullet”, *Computer* **20**(4), pp. 10–19 (April 1987).
- [18] Brooks, F.P. Jr., *The Mythical Man-Month: Essays on Software Engineering*, Second Edition, Addison Wesley, Reading, MA (1995).
- [19] Burgess-Yakemovic, K.C. and Conklin, J., “Report on Development Project Use of an Issue-Based Information System”, in *Proceedings of the ACM Conference on CSCW*, Los Angeles, CA (October 1990).
- [20] Carroll, J., Rosson, M.B., Chin, G., and Koenemann, J., “Requirements Development in Scenario-Based Design”, *IEEE Transactions on Software Engineering* **SE-24**(12), pp. 1156–1170 (December 1998).
- [21] CECOM, *Requirements Engineering and Rapid Prototyping Workshop Proceedings*, U.S. Army CECOM, Eatontown, NJ (1989).
- [22] Chan, W., Anderson, R.J., Beame, P., Burns, S., Modugno, F., Notkin, D., and Reese, J.D., “Model Checking Large Software Specifications”, *IEEE Transactions on Software Engineering* **SE-24**(7), pp. 498–520 (July 1998).
- [23] Dahl, O.-J., Dijkstra, E.W., and Hoare, C.A.R., *Structured Programming*, Academic Press, London, UK (1972).
- [24] Daugulis, A., “Time Aspects in Requirements Engineering: or ‘Every Cloud has a Silver Lining’”, *Requirements Engineering* **5**(3), pp. 137–143 (2000).
- [25] Dunsmore, A., Roger, M., and Wood, M., “Object-Oriented Inspection in the Face of Delocalization”, pp. 467–476 in *Proceedings of the Twenty-Second International Conference on Software Engineering (ICSE2000)*, Limerick, Ireland (June 2000).
- [26] Egyed, A., “A Scenario-Driven Approach to Traceability”, pp. 132–132 in *Proceedings of the Twenty-Third International Conference on Software Engineering (ICSE2001)*, Toronto, ON, Canada (June 2001).
- [27] Elspas, B., Levitt, K.N., Waldinger, R.J., and Waksman, A., “An Assessment of Techniques for Proving Program Correctness”, *Computing Surveys* **4**(2), pp. 81–96 (June 1972).
- [28] Elssamadisy, A. and Schalliol, G., “Recognizing and Responding to “Bad Smells” in Extreme Programming”, in *Proceedings of the Twenty-Fourth International Conference on Software Engineering (ICSE2002)*, Orlando, FL (May 2001).
- [29] Fagan, M.E., “Design and Code Inspections to Reduce Errors in Program Development”, *IBM Systems Journal* **15**(3), pp. 182–211 (1976).
- [30] Feldman, S.I., “Make — A Program for Maintaining Computer Programs”, *Software—Practice and Experience* **9**(4), pp. 224–265 (April 1979).

- [31] Forsberg, K. and Mooz, H., "System Engineering Overview", in *Software Requirements Engineering*, Second Edition, ed. R.H. Thayer and M. Dorfman, IEEE Computer Society Press, Washington (1997).
- [32] Gilb, T. and Graham, D., *Software Inspection*, Addison Wesley, Wokingham, UK (1993).
- [33] Glass, R.L., "The Realities of Software Technology Payoffs", *Communications of the ACM* **42**(2), pp. 74–79 (February 1999).
- [34] Goguen, J.A. and Tardo, J., "An Introduction to OBJ: A Language for Writing and Testing Formal Algebraic Program Specifications", in *Proceedings Conference on Specifications of Reliable Software*, Boston (1979).
- [35] Gotel, O.C.Z. and Finkelstein, A.C.W., "An Analysis of the Requirements Traceability Problem", pp. 94–101 in *Proceedings of the IEEE International Conference on Requirements Engineering*, IEEE Computer Society Press, Colorado Springs, CO (1994).
- [36] Griswold, W.G., Yuan, J.J., and Kato, Y., "Exploiting the Map Metaphor in a Tool for Software Evolution", pp. 265–274 in *Proceedings of the Twenty-Third International Conference on Software Engineering (ICSE2001)*, Toronto, ON, Canada (June 2001).
- [37] Hall, A., "Using Formal Methods to Develop an ATC Information System", *IEEE Software* **13**(66–76) (March 1996).
- [38] Heiss, J.J., "The Poetry of Programming", An Interview of Richard Gabriel, java.sun.com Report (3 December 2002), http://java.sun.com/features/2002/11/gabriel_qa.html.
- [39] Hester, S.D., Parnas, D.L., and Utter, D.F., "Using Documentation as a Software Design Medium", *The Bell System Technical Journal* **69**(8) (October 1981).
- [40] Highsmith, J. and Cockburn, A., "Agile Software Development: The Business of Innovation", *IEEE Computer* **34**(9), pp. 120–122 (September 2001).
- [41] Hoare, C.A.R., "An Axiomatic Basis for Computer Programming", *Communications of the ACM* **12**(10), pp. 576–580,585 (October 1969).
- [42] Jackson, M.A., "The Role of Architecture in Requirements Engineering", pp. 241 in *Proceedings of the IEEE International Conference on Requirements Engineering*, IEEE Computer Society Press, Colorado Springs, CO (April 18–22 1994).
- [43] Kemmerer, R.A., "Formal Verification of the UCLA Security Kernel: Abstract Model, Mapping Functions, Theorem Generation, and Proofs", Ph.D. Dissertation, Computer Science Department, UCLA (1979).
- [44] Kemmerer, R.A., "Testing Formal Specifications to Detect Design Errors", *IEEE Transactions on Software Engineering* **SE-11**(1), pp. 32–43 (January 1985).
- [45] Kiczales, G., Lamping, J., Mendhekar, A., Maeda, C., Lopes, C.V., Loingtier, J.-M., and Irwin, H., "Aspect-Oriented Programming", in *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, Springer, Finland (1997).
- [46] Knuth, D.E., *The Art of Computer Programming: Fundamental Algorithms*, Addison Wesley, Reading, MA (1969).
- [47] Knuth, D.E., *The Art of Computer Programming: Seminumerical Algorithms*, Addison Wesley, Reading, MA (1971).
- [48] Knuth, D.E., *The Art of Computer Programming: Sorting and Searching*, Addison Wesley, Reading, MA (1973).
- [49] Knuth, D.E., "Theory and Practice", *Theoretical Computer Science* **90**(1), pp. 1–15 (1991).
- [50] Lehman, M.M., "Programs, Life Cycles, and Laws of Software Evolution", *Proceedings of the IEEE* **68**(9), pp. 1060–1076 (September 1980).
- [51] Lehman, M.M. and Ramil, J.F., "The Impact of Feedback in the Global Software Process", *Journal of Systems and Software* **46**(2–3), pp. 123–134 (1999).
- [52] Leite, J.C.S.P. and Franco, A.P.M., "A Strategy for Conceptual Model Acquisition", pp. 243–246 in *Proceedings of the IEEE International Symposium on Requirements Engineering*, IEEE Computer Society Press, San Diego, CA (January 1993).
- [53] Lientz, B.P. and Swanson, E.B., "Characteristics of Application Software Maintenance", *Communications of the ACM* **21**(6), pp. 466–481 (June 1978).
- [54] Manna, Z. and Waldinger, R.J., "Toward Automatic Program Synthesis", *Communications of the ACM* **14**(3), pp. 151–165 (May 1971).
- [55] McConnell, S., "Daily Build and Smoke Test", *IEEE Software* **13**(4), pp. 144–143 (July/August 1996).
- [56] Müller, M.M. and Tichy, W.F., "Case Study: Extreme Programming in a University Environment", pp. 537–544 in *Proceedings of the Twenty-Third International Conference on Software Engineering (ICSE2001)*, Toronto, ON, Canada (June 2001).
- [57] Moriconi, M.S., "A Designer/Verifier's Assistant", *IEEE Transactions on Software Engineering* **SE-5**(4), pp. 387–401 (July 1979).
- [58] Myers, G.J., *The Art of Software Testing*, Wiley-Interscience, New York, NY (1979).
- [59] Naur, P. and Randell, B., "Software Engineering: Report on a Conference Sponsored by the NATO Science Commission", Garmisch, Germany, 7–11 October 1968, Scientific Affairs Division, NATO, Brussels, Belgium (January 1969).
- [60] Nuseibeh, B. and Easterbrook, S., "Requirements Engineering: A Roadmap", in *The Future of Software Engineering 2000*, ed. A. Finkelstein, ACM, Limerick, Ireland (June 2000).
- [61] Parnas, D.L., "On the Criteria to be Used in Decomposing Systems into Modules", *Communications of the ACM* **15**(2), pp. 1053–1058 (December 1972).
- [62] Parnas, D.L. and Clements, P.C., "A Rational Design Process: How and Why to Fake It", *IEEE Transactions on Software Engineering* **SE-12**(2), pp. 196–257 (February 1986).
- [63] Paulk, M.C., Curtis, B., Chrissis, M.B., and Weber, C.V., "Key Practices of the Capability Maturity Model", Technical Report, CMU/SEI-93-TR-25, Software Engineering Institute (February 1993).

- [64] Pressman, R., “Software According to Niccolò Machiavelli”, *IEEE Software* **12**(1), pp. 101–102 (January 1995).
- [65] Raymond, E., “Linux and Open-Source Success”, *IEEE Software* **16**(1), pp. 85–89 (January/February 1999).
- [66] Rochkind, M.J., “The Source Code Control System”, *IEEE Transactions on Software Engineering* **SE-1**(4), pp. 364–370 (December 1975).
- [67] Royce, W.W., “Managing the Development of Large Software Systems: Concepts and Techniques”, in *Proceedings of WesCon* (August 1970).
- [68] Schach, S.R., *Object-Oriented and Classical Software Engineering*, Fifth Edition, McGraw-Hill, New York, NY (2002).
- [69] Schwaber, K., “Agile Processes — Emergence of Essential Systems”, The Agile Alliance (2002), <http://www.agilealliance.org/articles/>.
- [70] Tichy, W., “RCS—A System for Version Control”, *Software—Practice and Experience* **15**(7), pp. 637–654 (July 1985).
- [71] Weidenhaupt, K., Pohl, K., Jarke, M., and Haumer, P., “Scenarios in System Development: Current Practice”, *IEEE Software* **15**(2), pp. 34–45 (1998).
- [72] Weinberg, G.M., *The Psychology of Computer Programming*, van Nostrand Reinhold, New York, NY (1971).
- [73] Wieringa, R., “Software Requirements Engineering: the Need for Systems Engineering and Literacy”, *Requirements Engineering Journal* **6**(2), pp. 132–134 (2001).
- [74] Wing, J.M., “A Specifier’s Introduction to Formal Methods”, *IEEE Computer* **23**(9), pp. 8–24 (September 1990).
- [75] Wirth, N., “Program Development by Stepwise Refinement”, *Communications of the ACM* **14**(4), pp. 221–227 (April 1971).