

SOFTWARE AND MIND

Andrei Sorin

EXTRACT

Chapter 6: *Software as Weapon*

**This extract includes the book's front matter
and chapter 6.**

Copyright © 2013, 2019 Andrei Sorin

**The free digital book and extracts are licensed under the
Creative Commons Attribution-NoDerivatives
International License 4.0.**

This chapter explains how the mechanistic fallacies lead to software delusions, and how the software elites use these delusions to exploit society.

The entire book, each chapter separately, and also selected sections, can be viewed and downloaded free at the book's website.

www.softwareandmind.com

SOFTWARE
AND
MIND

The Mechanistic Myth
and Its Consequences

Andrei Sorin

ANDSOR BOOKS

Copyright © 2013, 2019 Andrei Sorin
Published by Andsor Books, Toronto, Canada (www.andsorbooks.com)
First edition 2013. Revised 2019.

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, without the prior written permission of the publisher. However, excerpts totaling up to 300 words may be used for quotations or similar functions without specific permission.

The free digital book is a complete copy of the print book, and is licensed under the Creative Commons Attribution-NoDerivatives International License 4.0. You may download it and share it, but you may not distribute modified versions.

For disclaimers see pp. vii, xvi.

Designed and typeset by the author with text management software developed by the author and with Adobe FrameMaker 6.0. Printed and bound in the United States of America.

Acknowledgements

Excerpts from the works of Karl Popper: reprinted by permission of the University of Klagenfurt/Karl Popper Library.

Excerpts from *The Origins of Totalitarian Democracy* by J. L. Talmon: published by Secker & Warburg, reprinted by permission of The Random House Group Ltd.

Excerpts from *Nineteen Eighty-Four* by George Orwell: Copyright ©1949 George Orwell, reprinted by permission of Bill Hamilton as the Literary Executor of the Estate of the Late Sonia Brownell Orwell and Secker & Warburg Ltd.; Copyright ©1949 Harcourt, Inc. and renewed 1977 by Sonia Brownell Orwell, reprinted by permission of Houghton Mifflin Harcourt Publishing Company.

Excerpts from *The Collected Essays, Journalism and Letters of George Orwell*: Copyright ©1968 Sonia Brownell Orwell, reprinted by permission of Bill Hamilton as the Literary Executor of the Estate of the Late Sonia Brownell Orwell and Secker & Warburg Ltd.; Copyright ©1968 Sonia Brownell Orwell and renewed 1996 by Mark Hamilton, reprinted by permission of Houghton Mifflin Harcourt Publishing Company.

Excerpts from *Doublespeak* by William Lutz: Copyright ©1989 William Lutz, reprinted by permission of the author in care of the Jean V. Naggar Literary Agency.

Excerpts from *Four Essays on Liberty* by Isaiah Berlin: Copyright ©1969 Isaiah Berlin, reprinted by permission of Curtis Brown Group Ltd., London, on behalf of the Estate of Isaiah Berlin.

Library and Archives Canada Cataloguing in Publication

Sorin, Andrei

Software and mind : the mechanistic myth and its consequences / Andrei Sorin.

Includes index.

ISBN 978-0-9869389-0-0

1. Computers and civilization.
2. Computer software – Social aspects.
3. Computer software – Philosophy. I. Title.

QA76.9.C66S67 2013

303.48'34

C2012-906666-4

Don't you see that the whole aim of Newspeak is to narrow the range of thought?... Has it ever occurred to you ... that by the year 2050, at the very latest, not a single human being will be alive who could understand such a conversation as we are having now?

George Orwell, *Nineteen Eighty-Four*

Disclaimer

This book attacks the mechanistic myth, not persons. Myths, however, manifest themselves through the acts of persons, so it is impossible to discuss the mechanistic myth without also referring to the persons affected by it. Thus, all references to individuals, groups of individuals, corporations, institutions, or other organizations are intended solely as examples of mechanistic beliefs, ideas, claims, or practices. To repeat, they do not constitute an attack on those individuals or organizations, but on the mechanistic myth.

Except where supported with citations, the discussions in this book reflect the author's personal views, and the author does not claim or suggest that anyone else holds these views.

The arguments advanced in this book are founded, ultimately, on the principles of demarcation between science and pseudoscience developed by philosopher Karl Popper (as explained in "Popper's Principles of Demarcation" in chapter 3). In particular, the author maintains that theories which attempt to explain non-mechanistic phenomena mechanistically are pseudoscientific. Consequently, terms like "ignorance," "incompetence," "dishonesty," "fraud," "corruption," "charlatanism," and "irresponsibility," in reference to individuals, groups of individuals, corporations, institutions, or other organizations, are used in a precise, technical sense; namely, to indicate beliefs, ideas, claims, or practices that are mechanistic though applied to non-mechanistic phenomena, and hence pseudoscientific according to Popper's principles of demarcation. In other words, these derogatory terms are used solely in order to contrast our world to a hypothetical, ideal world, where the mechanistic myth and the pseudoscientific notions it engenders would not exist. The meaning of these terms, therefore, must not be confused with their informal meaning in general discourse, nor with their formal meaning in various moral, professional, or legal definitions. Moreover, the use of these terms expresses strictly the personal opinion of the author – an opinion based, as already stated, on the principles of demarcation.

This book aims to expose the corruptive effect of the mechanistic myth. This myth, especially as manifested through our software-related pursuits, is the greatest danger we are facing today. Thus, no criticism can be too strong. However, since we are all affected by it, a criticism of the myth may cast a negative light on many individuals and organizations who are practising it unwittingly. To them, the author wishes to apologize in advance.

Contents

	Preface	xiii
Introduction	Belief and Software	1
	Modern Myths	2
	The Mechanistic Myth	8
	The Software Myth	26
	Anthropology and Software	42
	Software Magic	42
	Software Power	57
Chapter 1	Mechanism and Mechanistic Delusions	68
	The Mechanistic Philosophy	68
	Reductionism and Atomism	73
	Simple Structures	90
	Complex Structures	96
	Abstraction and Reification	111
	Scientism	125
Chapter 2	The Mind	140
	Mind Mechanism	141
	Models of Mind	145

	Tacit Knowledge	155
	Creativity	170
	Replacing Minds with Software	188
Chapter 3	Pseudoscience	200
	The Problem of Pseudoscience	201
	Popper's Principles of Demarcation	206
	The New Pseudosciences	231
	The Mechanistic Roots	231
	Behaviourism	233
	Structuralism	240
	Universal Grammar	249
	Consequences	271
	Academic Corruption	271
	The Traditional Theories	275
	The Software Theories	284
Chapter 4	Language and Software	296
	The Common Fallacies	297
	The Search for the Perfect Language	304
	Wittgenstein and Software	326
	Software Structures	345
Chapter 5	Language as Weapon	366
	Mechanistic Communication	366
	The Practice of Deceit	369
	The Slogan "Technology"	383
	Orwell's Newspeak	396
Chapter 6	Software as Weapon	406
	A New Form of Domination	407
	The Risks of Software Dependence	407
	The Prevention of Expertise	411
	The Lure of Software Expedients	419
	Software Charlatanism	434
	The Delusion of High Levels	434
	The Delusion of Methodologies	456
	The Spread of Software Mechanism	469
Chapter 7	Software Engineering	478
	Introduction	478
	The Fallacy of Software Engineering	480
	Software Engineering as Pseudoscience	494

Structured Programming	501
The Theory	503
The Promise	515
The Contradictions	523
The First Delusion	536
The Second Delusion	538
The Third Delusion	548
The Fourth Delusion	566
The <i>GOTO</i> Delusion	586
The Legacy	611
Object-Oriented Programming	614
The Quest for Higher Levels	614
The Promise	616
The Theory	622
The Contradictions	626
The First Delusion	637
The Second Delusion	639
The Third Delusion	641
The Fourth Delusion	643
The Fifth Delusion	648
The Final Degradation	655
The Relational Database Model	662
The Promise	663
The Basic File Operations	672
The Lost Integration	687
The Theory	693
The Contradictions	707
The First Delusion	714
The Second Delusion	728
The Third Delusion	769
The Verdict	801
Chapter 8 From Mechanism to Totalitarianism	804
The End of Responsibility	804
Software Irresponsibility	804
Determinism versus Responsibility	809
Totalitarian Democracy	829
The Totalitarian Elites	829
Talmon's Model of Totalitarianism	834
Orwell's Model of Totalitarianism	844
Software Totalitarianism	852
Index	863

Preface

This revised version (currently available only in digital format) incorporates many small changes made in the six years since the book was published. It is also an opportunity to expand on an issue that was mentioned only briefly in the original preface.

Software and Mind is, in effect, several books in one, and its size reflects this. Most chapters could form the basis of individual volumes. Their topics, however, are closely related and cannot be properly explained if separated. They support each other and contribute together to the book's main argument.

For example, the use of simple and complex structures to model mechanistic and non-mechanistic phenomena is explained in chapter 1; Popper's principles of demarcation between science and pseudoscience are explained in chapter 3; and these notions are used together throughout the book to show how the attempts to represent non-mechanistic phenomena mechanistically end up as worthless, pseudoscientific theories. Similarly, the non-mechanistic capabilities of the mind are explained in chapter 2; the non-mechanistic nature of software is explained in chapter 4; and these notions are used in chapter 7 to show that software engineering is a futile attempt to replace human programming expertise with mechanistic theories.

A second reason for the book's size is the detailed analysis of the various topics. This is necessary because most topics are new: they involve either

entirely new concepts, or the interpretation of concepts in ways that contradict the accepted views. Thorough and rigorous arguments are essential if the reader is to appreciate the significance of these concepts. Moreover, the book addresses a broad audience, people with different backgrounds and interests; so a safe assumption is that each reader needs detailed explanations in at least some areas.

There is some deliberate repetitiveness in the book, which adds only a little to its size but may be objectionable to some readers. For each important concept introduced somewhere in the book, there are summaries later, in various discussions where that concept is applied. This helps to make the individual chapters, and even the individual sections, reasonably independent: while the book is intended to be read from the beginning, a reader can select almost any portion and still follow the discussion. In addition, the summaries are tailored for each occasion, and this further explains that concept, by presenting it from different perspectives.



The book's subtitle, *The Mechanistic Myth and Its Consequences*, captures its essence. This phrase is deliberately ambiguous: if read in conjunction with the title, it can be interpreted in two ways. In one interpretation, the mechanistic myth is the universal mechanistic belief of the last three centuries, and the consequences are today's software fallacies. In the second interpretation, the mechanistic myth is specifically today's mechanistic *software* myth, and the consequences are the fallacies *it* engenders. Thus, the first interpretation says that the past delusions have caused the current software delusions; and the second one says that the current software delusions are causing further delusions. Taken together, the two interpretations say that the mechanistic myth, with its current manifestation in the software myth, is fostering a process of continuous intellectual degradation – despite the great advances it made possible.

The book's epigraph, about Newspeak, will become clear when we discuss the similarity of language and software (see, for example, pp. 409–411).

Throughout the book, the software-related arguments are also supported with ideas from other disciplines – from the philosophies of science, of mind, and of language, in particular. These discussions are important, because they show that our software-related problems are similar, ultimately, to problems that have been studied for a long time in other domains. And the fact that the software theorists are ignoring this accumulated knowledge demonstrates their incompetence.

Chapter 7, on software engineering, is not just for programmers. Many parts

(the first three sections, and some of the subsections in each theory) discuss the software fallacies in general, and should be read by everyone. But even the more detailed discussions require no previous programming knowledge. The whole chapter, in fact, is not so much about programming as about the delusions that pervade our programming practices, and their long history. So this chapter can be seen as a special introduction to software and programming; namely, comparing their true nature with the pseudoscientific notions promoted by the software elite. This study can help both programmers and laymen to understand why the incompetence that characterizes this profession is an inevitable consequence of the mechanistic software ideology.

The book is divided into chapters, the chapters into sections, and some sections into subsections. These parts have titles, so I will refer to them here as *titled* parts. Since not all sections have subsections, the lowest-level titled part in a given place may be either a section or a subsection. This part is, usually, further divided into *numbered* parts. The table of contents shows the titled parts. The running heads show the current titled parts: on the right page the lowest-level part, on the left page the higher-level one (or the same as the right page if there is no higher level). Since there are more than two hundred numbered parts, it was impractical to include them in the table of contents. Also, contriving a short title for each one would have been more misleading than informative. Instead, the first sentence or two in a numbered part serve also as a hint of its subject, and hence as title.

Figures are numbered within chapters, but footnotes are numbered within the lowest-level titled parts. The reference in a footnote is shown in full only the first time it is mentioned within such a part. If mentioned more than once, in the subsequent footnotes it is abbreviated. For these abbreviations, then, the full reference can be found by searching the previous footnotes no further back than the beginning of the current titled part.

The statement “*italics added*” in a footnote indicates that the emphasis is only in the quotation. Nothing is stated in the footnote when the italics are present in the original text.

In an Internet reference, only the site’s main page is shown, even when the quoted text is from a secondary page. When undated, the quotations reflect the content of these pages in 2010 or later.

When referring to certain individuals (software theorists, for instance), the term “expert” is often used mockingly. This term, though, is also used in its normal sense, to denote the possession of true expertise. The context makes it clear which sense is meant.

The term “elite” is used to describe a body of companies, organizations, and individuals (for example, the software elite). The plural, “elites,” is used when referring to several entities within such a body.

The issues discussed in this book concern all humanity. Thus, terms like “we” and “our society” (used when discussing such topics as programming incompetence, corruption of the elites, and drift toward totalitarianism) do not refer to a particular nation, but to the whole world.

Some discussions in this book may be interpreted as professional advice on programming and software use. While the ideas advanced in these discussions derive from many years of practice and from extensive research, and represent in the author’s view the best way to program and use computers, readers must remember that they assume all responsibility if deciding to follow these ideas. In particular, to apply these ideas they may need the kind of knowledge that, in our mechanistic culture, few programmers and software users possess. Therefore, the author and the publisher disclaim any liability for risks or losses, personal, financial, or other, incurred directly or indirectly in connection with, or as a consequence of, applying the ideas discussed in this book.

The pronouns “he,” “his,” “him,” and “himself,” when referring to a gender-neutral word, are used in this book in their universal, gender-neutral sense. (Example: “If an individual restricts himself to mechanistic knowledge, his performance cannot advance past the level of a novice.”) This usage, then, aims solely to simplify the language. Since their antecedent is gender-neutral (“everyone,” “person,” “programmer,” “scientist,” “manager,” etc.), the neutral sense of the pronouns is established grammatically, and there is no need for awkward phrases like “he or she.” Such phrases are used in this book only when the neutrality or the universality needs to be emphasized.

It is impossible, in a book discussing many new and perhaps difficult concepts, to anticipate all the problems that readers may face when studying these concepts. So the issues that require further discussion will be addressed online, at www.softwareandmind.com. In addition, I plan to publish there material that could not be included in the book, as well as new ideas that may emerge in the future. Finally, in order to complement the arguments about traditional programming found in the book, I have published, in source form, some of the software I developed over the years. The website, then, must be seen as an extension to the book: any idea, claim, or explanation that must be clarified or enhanced will be discussed there.

Software as Weapon

Our discussion in the previous chapter – the use of language to deceive and exploit – had a dual purpose. First, we benefited by discovering the methods employed by charlatans to turn language from a means of communication into a tool of manipulation. From the present-day deceivers to Orwell’s totalitarian society, we saw that the aim of linguistic manipulation is to distort knowledge. And this is accomplished by forcing us to commit the mechanistic fallacies of reification and abstraction. Specifically, the charlatans force us to create isolated and impoverished knowledge structures in our minds. They prevent us from using those natural capabilities of the mind that allow us to process complex structures, and our mental processes are reduced to the level of machines.

But the main reason for the study of linguistic manipulation was to demonstrate the immense power inherent in language – the power to control minds. Then, if we understand the similarity of language and software, and of linguistic manipulation and software manipulation, we can better appreciate how the same power can be attained through software.

We already know that language and software fulfil a similar function: mirroring the world and communicating with it. And we also know that they work in a similar manner: by creating systems of interacting structures from linguistic and software entities, respectively. Having established that reification

and abstraction are the methods whereby language is used to control minds, what we must do now is study our *software* delusions, and how the software elites are using these delusions to exploit us.

What we will find is that the very methods we recognize as means of deception and manipulation in language – forcing us to separate knowledge structures and to start from higher levels of abstraction – form in software an open doctrine. Whereas the language charlatans try to *hide* the fact that they are employing these methods, the software charlatans are quite candid about it: separating software structures and starting from higher levels, they keep telling us, is the most effective way to develop software.

But software is becoming as important as language, so the conclusion must be that our elites are attaining through software the same power over human minds that the elite in Orwell's hypothetical society attains through language. Thus, while no elite in a real society can ever attain the power to control *language* to the degree that this is done in Orwell's society, our software elites already have that power in controlling *software*. And, if not stopped, they will eventually control our minds as effectively as the elite controls minds in Orwell's society.

A New Form of Domination

The Risks of Software Dependence

I have stated that the elites can control knowledge by means of software just as they can by means of language, but I must clarify this point. I am not referring here to the *direct* use of software to control, acquire, or restrict knowledge. This discussion is *not* concerned with such well-known software dangers as allowing an authority to decide what information is to be stored in databases, or allowing the centralized collection of information about individuals. Nor is it concerned with the use of deceptive language in software propaganda, as in calling everything “technology,” or “solution.” Important as these dangers are, they are insignificant compared to the dangers we face when an elite controls *the way we create and use software*.

The other dangers we understand, and if we understand them we can perhaps deal with them. But we have yet to understand what it means for a society to depend on software as much as it depends on language; and consequently, we do not realize that it is just as important to prevent an elite

from controlling software as it is to prevent one from controlling language. This ignorance can be seen in the irresponsible attitudes of our political leaders, of our corporations and educational institutions, and ultimately of every one of us: we are watching passively as the software elites are increasing their power and control year after year; and we continue to trust and respect them, even as they are creating a world where the only thing left for us to do is to operate their devices.

Thus, while the other forms of software abuse would lead to familiar forms of exploitation, what concerns us here is *a new form of domination*. We are facing a new phenomenon, a new way to control knowledge and thought, which could not exist before we had computers. The dependence of a society on software is a new phenomenon because software and programming are new phenomena. We have been inventing tools for millennia, but the computer is unique in that it is programmable to a far greater extent than any other tool we have had. Software, therefore, is what gives the computer its potency; and the act of programming is what controls this potency. No other human activity – save the use of language – is as far-reaching as programming, because no other activity involves something as potent as software.

This is the first time since humans developed languages that we have invented something comparable in scope or versatility. Software resembles language more than anything else: both systems permit us to mirror the world in our minds and to communicate with it. At the same time, software is sufficiently different from language to mask the similarity (we can easily invent new programming languages, for example, but not natural ones). As a result, we fail to appreciate the real impact that software has on society, and the need for programming expertise and programming freedom. And we have fallen victim to the fallacies of the software myth: the belief that software is a kind of product, and that software applications must be built as we build appliances; the belief that we need elaborate tools for these manufacturing projects, and hence a software industry to supply these tools; the belief that development methodologies and environments can be a substitute for programming expertise; the belief that it is better to program and maintain complex business systems by employing large teams of inexperienced programmers, analysts, and managers, instead of one professional; and the belief that the highest programming skills that human minds can attain, and that society needs, are those possessed by the current practitioners.



It is not surprising that we are unprepared for the consequences of programming, since we did not take sufficient time to learn what programming really is.

Originally, we set out merely trying to develop a particular kind of machine – a fast, programmable calculator. Instead, we stumbled upon a system that gives us a whole new way to use our minds, to communicate, to represent the world. But we continue to regard programming as we did originally, as an extension to the engineering effort required to build the hardware; that is, as an activity akin to manufacturing, and which must be performed in the same fashion. We still fail to see that the skills needed to program computers are more akin to those needed to use language. Programming projects cannot be neatly broken down, like manufacturing activities, into simpler and simpler tasks. Programming skills, therefore, entail a capacity for complex structures. They can be acquired only through lengthy exposure to the phenomena arising from developing and maintaining large and complex applications.

There can be little doubt that within a few decades humans will interact with the world around them by means of software as much as they do now by means of language. Software lends itself to this task just like language, and there is no reason why we should not depend on our newly discovered programming capabilities, just as we depend on our linguistic capabilities. We must ensure, however, the right conditions: first, programmers must have the same competence with software as normal humans have now with language; and second, the activities involving programming and software must be, like those involving language, free from controls or restrictions. A society that allows an elite to control its software and programming faces the same danger as a society that allows its *language* to be controlled: through language or through software, the elite will eventually control all knowledge and thought, and will reduce human minds to the level of machines.

The form of domination that we are studying here can emerge, therefore, when a society depends on software but lacks the necessary programming competence. As Orwell points out, an elite could achieve complete control through language only by forcing us to replace our language with an impoverished one, like Newspeak, which demands no intelligence. In our present-day society, exploitation by way of language is necessarily limited, because we are all competent language users. Few programmers, however, attain a level of software competence comparable to our linguistic competence. Accordingly, the world of programming is already an Orwellian world: it resembles a society that lacks linguistic competence.

Our dependence on programming aids, and on the organizations behind them, stems from the incompetence of programmers. Programmers need these devices because they are not permitted to attain the level of programming competence of which human minds are naturally capable. But programming aids are only poor substitutes for programming expertise, because, unlike minds, they can only deal separately with the various aspects of programming.

As a result, applications based on these substitutes cannot represent the world accurately, just as statements in Newspeak cannot.

And, just as people restricted to Newspeak cannot realize how limited their knowledge is, *we* cannot realize how limited our *programming* knowledge is, because this is the only kind of programming we have. Just as the people in Orwell's society are forced to depend on the linguistic tools provided by their elite, and their knowledge is shaped and restricted by these tools, our programmers are forced to depend on the devices provided by the software companies, and their knowledge is similarly shaped and restricted. Only mechanistic software concepts, only beliefs that reinforce the software myth, can enter their minds. Programming expertise for them means expertise in the use of substitutes for programming expertise.

By preventing programming competence, then, an elite can use software to control and exploit society, just as language could be used if we lacked *linguistic* competence. The programming aids, and the resulting applications, form a complex world that parallels the real world but has little to do with it. Their chief purpose is to support a large software bureaucracy, and to prevent the emergence of a body of competent and responsible programmers. And if the software bureaucrats no longer deal with the real world, we have to reshape our own views to match theirs. To the extent that our society depends on software, and hence on this software bureaucracy, we *all* live in an Orwellian world: we are all forced to perceive our work, our values, our expectations, our responsibilities, in ways that serve the interests of the software elites.

It is unlikely that Orwell's extreme form of mind control through language can ever happen in the real world, but this is unimportant. Orwell's world is a model, not a prophesy. We must appreciate its value as model, therefore, rather than feel self-complacent because it cannot happen. And when we study it, we recognize that its importance is growing as our dependence on software is growing, because this dependence increases the possibility of an elite controlling our minds through software as the elite in the model does through language.

In our current software culture, the degree of control that an elite can attain through software is not limited by an existing condition, as control through language is limited by our linguistic competence; it rests solely on how much we depend on software. The reason the software elites do not have complete control over our minds today is not our software competence, but the fact that we do not yet depend completely on software. And if our dependence on software is growing, by the time we depend on software as much as we depend now on language it will be too late to do anything. At that point, to use Orwell's model, we will live in a world where Newspeak finally replaces English. Our chief means of thinking, of communicating, of representing the

world, will be a simple system requiring only mechanistic knowledge – not because software structures cannot involve complex knowledge, but because there will be no one to create or use the kind of software that requires the full capacity of the mind.

Dependence on software, coupled with software ignorance and programming incompetence – this is what the software elites are trying to achieve. They are persuading us to give up our dependence on knowledge and skills (means through which we *know* how to become competent) and to replace it with a dependence on software (means which they control, and through which they can *prevent* us from becoming competent).

The Prevention of Expertise

1

We probably fail to recognize software domination because the idea of mind control through software is so incredible. Before we had software, only political organizations could carry out such a totalitarian project. And we have yet to accept the fact that an elite can control society through software as effectively as a political elite could through traditional means of domination.

To understand this danger, we must make the most of what we know today. We cannot afford merely to wait and see, because the resulting conditions would likely be irreversible. We must study, for example, the similarity between the role of software in society and that of language. Since we all agree that language can be used to control and restrict thought, we must ensure complete software freedom even if we still cannot see clearly how the software elites can turn software into a means of domination. We should simply assume that they will use software as they would language, had they the opportunity to control language as they do software.

Even more importantly, we must study those aspects of society that are *already* controlled by the software elites: those aspects that form the world of programming itself. Studying the world of programming affords us a glimpse of the future, of the time when the entire society will be controlled by these elites. It was easy to degrade the notion of programming expertise because, this being a new field, there were no established values, as there are in the traditional professions. As a result, the highest level of expertise we believe to be needed in programming is one that in other professions would be considered the level of novices.



We have been involved with software for more than half a century, so by now we could have had a sufficient number of expert programmers; namely, men and women whose skills represent *the utmost that human minds can attain in the domain of programming*. This is how we define expertise in other professions, and this is what we should expect of programmers. Instead, what we find is a software *bureaucracy*: a social system whose chief doctrine is the *prevention* of programming expertise.

We have programmers who are incapable of performing anything but small and isolated programming tasks, and who are not even expected to do more. We have managers who read “success stories” in childish computer publications, and search for ready-made applications and other programming substitutes instead of allowing their own programmers to gain the necessary skills. We have professors and gurus who teach the principles of “software engineering” – which claim that programming is like manufacturing, so what we need is unskilled labourers who know only how to assemble “prefabricated software components.” We have software companies that supply practitioners with an endless series of “software tools” – elaborate development and business systems that promise to eliminate the need for programming. And, addressing the incompetence engendered by this corrupt culture, there are thousands of books, magazines, newspapers, brochures, advertisements, catalogues, trade shows, newsletters, courses, seminars, and online sources, all offering “solutions.”

Few people realize that this whole bureaucracy could be replaced with a relatively small number of real, expert programmers. This is true because only a fraction of the work performed by the current practitioners is actually useful; that is, directly related to the creation and maintenance of applications. Most of their work consists in solving the problems generated by their dependence on aids and substitutes.

We have no equivalent bureaucracy in other professions. We have surgeons, pilots, engineers, musicians, military commanders, writers, repairmen, and so forth. And we understand that, to reach expertise in a difficult profession, an individual needs many years of education, training, and practice, a sense of personal responsibility, and perhaps special talents as well. We don't attempt to replace a surgeon with a dozen ignorant individuals, and defend the decision by claiming that the work of a surgeon can be broken down into simpler tasks, as in manufacturing.

We don't do this in other professions because we took the time to determine what is the *highest* level that human beings can attain in those fields. We made *that* level our definition of expertise, and we measure everyone's performance against that level. We understand that the more difficult the profession, the longer it takes to attain expertise, and the fewer the individuals who can

succeed; and we give these individuals the time and the opportunity to develop their capabilities. We never tried to contend with this problem by reducing our expectations, as we do in programming. We never concluded that, given the urgent need for surgeons, the answer is to debase the definition of expertise to match the level of the available, of the inexperienced, individuals.

We treat programming differently from other professions because this serves the interests of the software elites. In just a few years, an unprecedented propaganda system has made the software myth the greatest system of belief in history, and we now take for granted in the domain of programming, notions that we would dismiss as absurd in other domains. The software practitioners have become a powerful, self-serving bureaucracy, but we continue to regard them as saviours. The reason we fail to see that they are *exploiting* society, not serving it, is that we have no other form of programming as measure. Programming controlled by a bureaucracy is the only programming we know, the only kind we have ever had.

An important element of the software myth is the belief that the typical work currently performed by programmers represents the highest level of competence we should expect in this profession. And if they have difficulty with their applications, it is not greater programming knowledge that they need, but more programming aids and substitutes. Thus, individuals with just a year or two of training and experience – which consist largely in the use of aids and substitutes – have reached the highest knowledge expected of them. The only knowledge they will acquire from then on is how to use the *future* aids and substitutes. This doctrine fits well within the ideology promoted by the software elites, as it ensures continued incompetence among programmers. It also ensures the complete dependence of programmers, and of those using their applications, on the software companies supplying the aids and substitutes. Lastly, this doctrine serves to weaken the programmers' sense of responsibility: what they perceive as their main concern is the problems generated by the aids and substitutes, rather than the real social or business problems that software is supposed to solve.

As a result, no matter how many years of practice programmers have behind them, their real programming experience remains as it was after the first year or two. This is true because the aids and substitutes limit their work to simple and isolated bits of programming, whereas successful application development demands the capacity to deal with many software processes simultaneously. This incompetence also explains why most applications are inadequate, and why most programming work consists in replacing existing applications, which programmers cannot keep up to date.



If we can benefit from studying the similarity of software and language, and from studying the world of programming and the delusions of the software myth, then we can benefit even more from studying these topics together; specifically, from studying the link between the mechanistic ideology and the incompetence of programmers.

We should regard the world of programming as the result of an unintended social experiment: an attempt to replace human expertise with software. The experiment has failed, but we can learn a great deal from this failure. We must create, to begin with, the social and business environment where a body of expert programmers can evolve. The software elites are doing everything in their power to prevent this, of course, since widespread programming incompetence is a critical factor in their plan of domination. A true programming profession will not only stop the flow of trillions of dollars from society to the software elites and bureaucrats, but will lead to better software and, ultimately, greater benefits from computers.

Moreover, by abolishing the software bureaucracy we will prevent the software elites from corrupting other aspects of society. For they are using the power gained from controlling the world of software, to degrade other professions and occupations just as they have degraded programming. If allowed to continue, they will soon force us all to depend on knowledge substitutes instead of our minds. Like programmers, we will all be reduced to the level of novices. As programmers do now, we will all live in a world where expertise is neither possible nor necessary, where the only thing left to do is to operate the devices supplied by the software elites.

2

To understand the concept of software domination, we must start by recalling what we learned in chapter 2 about the mind. We can acquire the most diverse kinds of knowledge and skills: using language, recognizing faces, playing games, programming computers – the kind of knowledge we all share simply by living in a society, as well as specialized knowledge related to individual lifestyles and occupations. But all knowledge and skills, ultimately, involve our mind's capacity for complex structures. When exposed to a new phenomenon, and hence to the new knowledge embodied in that phenomenon, we start by noting the *simple* structures that make it up, the patterns and regularities. What we note, in other words, is those aspects that can be represented with facts, rules, and methods. Being limited to simple structures, our performance at this point resembles that of a software device. We progress from novice to expert by being exposed to that phenomenon repeatedly. This permits our

mind to discover, not only more structures, but also the *interactions* between structures; and this in turn permits it to create a replica of the *complex* structures that make up the phenomenon. Thus, when we reach expertise our performance exceeds that of software devices, which are forever restricted to simple structures.

We communicate with the world around us through our senses, which receive information in the form of simple structures (patterns of symbols and sounds, for instance). Complex structures, therefore, can exist only in the phenomenon itself and in the mind; we cannot acquire them through our senses directly from the phenomenon, or from another mind. Consequently, the only way to attain expertise in a given domain is by giving our mind the opportunity to create the complex structures which reflect the phenomena of that domain. And this the mind can do only through repeated exposure to those phenomena; in other words, through personal experience.

Human acts require the capacity for complex structures because most phenomena we face consist of interacting structures. They consist of entities (objects, persons, processes, events) that have many attributes, and belong therefore to many structures at the same time – one structure for each attribute. To put this differently, we can always view ourselves and our environment from different perspectives, while the entities that constitute our existence are the same. So the entities form many structures; but the structures interact, because they share these entities. We rarely find a structure – a particular aspect of our life – whose links to the other structures are so weak that it can be extracted from the complex whole without distorting it or the others.

This recapitulation was necessary in order to remind ourselves of the conclusion reached in chapter 2, and its significance. We note that most mental processes, most knowledge and skills, involve *complex* structures. And we note also that software devices are based on *simple* structures. As substitutes for human intelligence, therefore, software devices are useful only for the rare tasks that can be represented with simple structures; specifically, those tasks that can be separated from others.

On the one hand, then, practically everything we do involves the full capacity of the mind, and cannot be broken down into simpler mental processes. On the other hand, we agree to depend more and more, in almost every domain, on software devices – which attempt to eliminate the need for expertise by reducing knowledge to simple structures. How can we explain this contradiction?

Our software delusions stem from our mechanistic delusions. Our most popular theories of mind claim that human intelligence can be represented with mechanistic models – models based on precise diagrams, rules, methods, and formulas. And, even though these theories keep failing, we also believe

now that it is possible to represent intelligence with mechanistic *software* models. Thus, the promoters of mind mechanism can claim, for the first time, to have *actual* devices – software devices – that emulate human intelligence. Anyone with a computer can now perform any task, including tasks requiring knowledge that he lacks. All he needs to do is purchase a software device which contains that knowledge.

We have always used tools to simplify tasks, or to improve our performance; so the idea that a device can enhance certain types of knowledge and skills, or help us perform some tasks faster or better, is not new. If we view software as a device of this kind, the claims are easily justified: the computer, with the programs that run on it, is the most versatile tool we have ever invented; and it can enhance our capabilities in many tasks.

The software claims, though, do not stop at the kind of claims traditionally advanced for devices. The claims are extended to encompass *intelligent* acts; that is, acts involving non-mechanistic knowledge, and hence complex knowledge structures. But devices can represent only *simple* structures. So, to replace those acts with software, we must first separate the complex knowledge structure into several simple ones.

Software domination, thus, starts when we are tempted to commit the fallacy of reification. We believe the claim that knowledge and skills can be replaced with software devices because we already believe that intelligent acts can be broken down into simpler intelligent acts. This belief tempts us to reify the phenomenon of intelligence, and commit therefore, with software, the fallacy already committed by the mechanistic theories of mind: the separated knowledge structures are no longer what they were when part of the complex knowledge; they lose the interactions, so even when we manage to represent them faithfully with software, the knowledge embodied in them is not the same as the original, complex knowledge.

But reification is only the first step. Now that we have independent structures, we are tempted to start from higher levels of abstraction within each structure as we replace it with software. We can be more productive, the experts tell us, if we avoid the low levels of software and start with higher-level elements – with elements that already contain the lower levels. Thus, we also commit the second fallacy, abstraction: we believe that we can accomplish the same tasks as when starting with low-level elements. Starting from higher levels impoverishes the structure by reducing the number of alternatives for the value of the top element; that is, the top element of a software structure that is already a reified, and hence inaccurate, representation of the real knowledge. What this means in practice is that an inexperienced person will accomplish by means of software devices only a fraction of what an experienced person will with his mind alone.

The two fallacies can be seen clearly in the domain of programming. We are told that the most effective way to develop applications is by starting from high levels of abstraction. Specifically, we should avoid programming as much as possible, and use instead software entities that already include many elements: ready-made applications (or, at least, ready-made modules and components), and the built-in operations provided by development tools and environments. To benefit from these high levels, however, we must view our applications, mistakenly, as separable software processes. Each business or software practice, each case of shared data or operations, is a process; and each process represents one aspect of the application, one structure (see “Software Structures” in chapter 4). These structures exist at the same time and use the same software entities, so it is their totality that constitutes the application. If we separate them, we may indeed manage to program each structure starting from higher-level elements; but the resulting application will reflect only poorly the original requirements. We *can* create applications that mirror reality, but only if we have the expertise to start from low levels and to deal with all the processes together.

3

We are now in a position to understand the concept of software domination. The software elites are exploiting our mechanistic delusions; specifically, our belief that software can be a substitute for non-mechanistic knowledge. We see software successfully replacing human minds in *some* tasks, and we trust the elites when they promise us similar success in other tasks. We believe that there is only a quantitative, not a qualitative, difference between the knowledge involved in mechanistic and non-mechanistic tasks, and we allow the elites to exploit this belief.

So the mechanistic delusions act like a trap. If we believe that a given task can be replaced with software, we do not hesitate to depend on that software and on the organization behind it. We enthusiastically devote our time to that software, instead of using it to gain knowledge. But if the task is non-mechanistic, the time we devote to software will be wasted. By luring us with the promise of immediate answers to our complex problems, the elites prevent us from developing our minds. So, not only does software fail to solve those problems, but it also prevents us from gaining the knowledge whereby we *could* solve them. In the end, we have nothing – neither a useful software expedient, nor the knowledge needed to perform the task on our own. We are caught, thus, in the software variant of the traditional mechanistic trap: we believe in the existence of mechanistic solutions for non-mechanistic problems, so we restrict ourselves to precisely those methods that cannot work.

We get to depend on software because the promise is so enticing. The promise is, essentially, that software can solve important problems in a particular field by acting as a substitute for the knowledge needed in that field. So, instead of taking the time to acquire that knowledge, we can solve the problems right away, simply by buying and operating a piece of software. And the reason we believe this promise is that we see similar promises being fulfilled in tasks like calculations and data processing, where software does indeed allow novices to display the same performance as experts. For tasks involving interacting knowledge structures, however, the promise cannot be met. Software can be a substitute only for knowledge that can be neatly broken down into smaller, simpler, and independent pieces of knowledge – pieces which in their turn can be broken down into even smaller ones, and so on – because this is the only way to represent knowledge with software.

When we depend on software in situations involving complex knowledge structures, what we notice is that it does not work as we expected. But, while software fails to provide an answer, we never question the mechanistic assumptions. We recognize its inadequacy, but we continue to depend on it. And this dependence creates a new *kind* of problems: software-related problems. We don't mind these new problems, though, and we gladly tackle them, because we think that by solving *them* we will get the software to solve our original problems. Most software problems involve isolated knowledge structures, and have therefore fairly simple, mechanistic solutions. Since these solutions often entail additional software, they generate in their turn software problems, which have their own solutions, and so on.

The more software problems and solutions we have, the more pleased we are with our new preoccupations: all the time we seem to be *solving* problems. So we spend more and more time with these software problems, when in fact they are spurious problems generated by spurious preoccupations, in a process that feeds on itself. We interpret our solutions to software problems as progress toward a solution to our real problems, failing to see that, no matter how successful we are in solving software problems, if the original problems involve complex knowledge we will never solve *them* with software. In the end, content with our software preoccupations, we may forget the real problems altogether.



Bear in mind that it is not software dependence in itself that is dangerous, but the combination of this dependence with software ignorance. There is nothing wrong in depending on software when we possess software expertise. It is only when we separate software structures from various knowledge structures, and

when we restrict ourselves to high-level elements, that our dependence on software can be exploited; in other words, when our software knowledge is at a mechanistic level.

Recall the language analogy. We attain linguistic competence by starting with *low-level* linguistic elements – morphemes and words, with all their uses and meanings – and by creating language structures that interact with various knowledge structures present in the mind. This is how we form the complex mental structures recognized as intelligence. No linguistic competence or intelligence would be possible if we had to create our language structures starting with ready-made sentences and ideas, or if we treated them as independent structures. Similarly, software expertise can be attained only by starting with *low-level* software elements, and by treating the software structures that we create with our programs, not as independent structures, but as part of the complex structures that make up our affairs. And this is precisely what the software elites are *preventing* us from doing.

We are exploited through software because of the belief that it is possible to benefit from software without having to develop software expertise. We understand why we could not benefit from *language* without having linguistic competence, but we fail to see this for software. When we trust the elites and get to create and use software in the manner dictated by them, we end up with a combination of several weaknesses: programming incompetence; failure to solve our problems with software, because of the inadequacy of our applications; a growing preoccupation with spurious, software-related problems; and a perpetual dependence on the elites for solutions to both the real and the spurious problems – a dependence that is futile in any case, because only we, with our minds, can hope to accomplish those tasks that require complex knowledge. It is not difficult for the elites, then, to exploit us.

The Lure of Software Expedients

1

Software domination is based on a simple stratagem: consuming people's time. Forcing us to waste our time is the simplest way to keep us ignorant, since, if we spend this time with worthless activities, we cannot use it to improve our minds. The software elites strive to keep us ignorant because only if ignorant can normal people be turned into bureaucrats, into automatons.

Our time, in the end, is all we have, our only asset. Whether we count the hours available in a day or the years that constitute a life, our time is limited. What we do with our time, hour by hour, determines what we make of our

lives. We can squander this time on unimportant pursuits, or use it to expand as much as we can our knowledge and skills. In the one case we will accomplish whatever can be done with limited knowledge, and we will probably live a dull life; in the other case we will make the most of our minds, and we have a good chance to live a rich life and to make a contribution to society.

It is not too much to say that, as individuals living in a free society, we are, each one of us, responsible for a human life – our own – and we have an *obligation* to make the most of it. Like freedom itself, realizing our human potential is a *right* we all have; but, just like freedom, this right is also a *duty*, in that we are all responsible for its preservation by defending it against those who want to destroy it. Specifically, we must strive to expand our minds *despite* the attempts made by an elite to keep us ignorant. Only thus, only when each individual *and* each mind counts, can the idea of freedom survive.

Conversely, preventing an individual from realizing his or her potential, from making the most of his or her mind, amounts in effect to an attempt to destroy a life; so it must be considered a crime nearly as odious as murder. Seen from this perspective, our software elites could be described as criminal organizations, since forcing us to squander our time is one of the principles of their ideology.

The software elites consume our time by creating an environment where our activities are far below our natural mental capabilities. When we depend on their concepts and devices, we end up spending most of our time acquiring isolated bits of knowledge, or solving isolated problems. Being simple and mechanistic, these activities do not allow us to create complex knowledge structures in our minds – the kind of knowledge that constitutes skills and experience. We can recognize this in that our capabilities do not progress on a scale from novice to expert, as they do in the traditional fields of knowledge. (Thus, no matter how many of these problems we solve, the next one will demand about as much time and effort.) In any case, mechanistic concepts cannot help us to solve our *complex* problems; so we are wasting our time both when acquiring the mechanistic knowledge and when using it.

Without exception, the software devices are presented as simple, easy to use, requiring little knowledge, and demanding an investment of just a few minutes, or perhaps a few hours. It is a sign of our collective naivety that, in a world which is becoming more complex by the day, we believe in the existence of some devices that can provide immediate answers to our problems, and that this power is ours to enjoy just for the trouble of learning how to operate them. This childish belief can be understood only by recognizing it as the software variant of our eternal craving for salvation: the performance of some simple acts, we think, will invoke the assistance of fabulous powers. This is the same belief that permitted so many other elites to exploit us in the past.

Many of these devices are indeed simple, just as their promoters claim. But they are simple precisely because they are generally useless, because they *cannot* solve our complex problems. As we saw, they are based on the delusion that complex structures – our problems, and the knowledge required to solve them – can be broken down into simple structures, and hence replaced with mechanistic expedients like software devices. The mechanistic knowledge required to use the software, together with the mechanistic capabilities of the software itself, is believed to provide a practical substitute for the complex knowledge required to solve the problems. But no combination of simple structures can replace a complex knowledge structure, because this kind of knowledge can develop only in a mind, and only after much learning and practice.



When we trust the software elites, we are exploited more severely and in more ways than we think. For, the waste of time and the prevention of knowledge reinforce each other, impoverishing our lives and degrading our minds to the point where we can no longer understand what is happening to us. This is an important point, and we must dwell on it.

We saw that the promise of immediate solutions to our problems tempts us to depend on mechanistic software expedients – ready-made pieces of software, or software tools and aids. Complex knowledge becomes both unnecessary and impossible, as all we need to do is select and combine various operations within the range of alternatives permitted by a device. We agree to depend on mechanistic expedients because we believe this is the most effective way to accomplish a given task, when in fact most tasks require *non-mechanistic* knowledge.

But it is important to note that there is a second process at work here: if we spend our time engaged in activities requiring only mechanistic thinking, we lose the *opportunity* to develop complex, non-mechanistic knowledge. We have the *capacity* for non-mechanistic knowledge, but we will not take the time to *develop* it as long as we believe that the simpler, mechanistic knowledge suffices. So our knowledge remains at mechanistic levels, and we continue to depend on the mechanistic expedients, even though they are inferior to our own capabilities.

What distinguishes human minds from devices is the capacity to develop complex knowledge. But, as we saw in chapter 2, the only way to attain complex knowledge is through personal experience: by engaging in activities that require that knowledge, by being exposed repeatedly to the complex phenomena that embody it. Thus, in some domains we need many years of

learning and practice to attain the knowledge levels recognized as expertise. Mechanistic knowledge, on the other hand, can be acquired fairly quickly. (See pp. 155–157.)

The promise of mechanistic solutions to complex problems is, then, a trap. When inexperienced, and hence limited to mechanistic knowledge, the mechanistic expedients do indeed exceed our skills. It is only later, when we develop non-mechanistic knowledge, that we will outperform them. But we will never reach that level if we get to depend on mechanistic expedients from the start, because this very dependence deprives us of the opportunity to develop non-mechanistic knowledge.

This degradation – restricting us to mechanistic thinking, to a fraction of our mental capabilities – is the goal of the software elites when tempting us with mechanistic expedients. The prevention of non-mechanistic knowledge is a critical element in their plan of domination, because they must ensure that we remain inferior to their devices, and hence dependent on them.

We can choose only one of the two alternatives: either pursue activities demanding mechanistic knowledge (because they are easy and immediately accessible), or take the time to develop non-mechanistic knowledge. Mechanistic knowledge (following rules and methods, operating a software device) we can quickly acquire at any time, while non-mechanistic knowledge (the experience to perform complex tasks, the creativity to solve important problems) requires many years of learning and practice.

The software elites encourage us to choose the first alternative. This choice brings immediate rewards and is hard to resist, but it restricts us forever to mechanistic thinking. To prefer the second alternative, we must appreciate the potential of our minds. This choice amounts, in effect, to an investment in ourselves: we decide to forgo some easy and immediate benefits, and, instead, take the time to develop our minds. But we can make this choice only if we already have an appreciation of non-mechanistic knowledge, only if we realize how much more we can accomplish later, when we attain this type of knowledge. And we can develop this appreciation only if, when young or when novices in a particular field, we note around us both mechanistic and non-mechanistic knowledge, and learn to respect those who possess the latter – because their skills exceed ours by far.

The software elites, however, are creating a culture that fosters mechanistic thinking – a culture where non-mechanistic capabilities offer no benefits, as we are all expected to stay at about the same skill level. More and more, in one occupation after another, the only thing we have to know is how to use a software system. The notions of expertise, creativity, professionalism, and responsibility are being degraded to mean simply the skill of following methods and operating devices. As we depend increasingly on mechanistic

knowledge alone, non-mechanistic knowledge is becoming redundant: we have fewer and fewer opportunities to either develop it or use it.

By creating a culture where all we need is mechanistic knowledge, the elites make it impossible for us to discover the superiority of *non-mechanistic* knowledge. We are trapped in a vicious circle: we start by being inexperienced and hence limited to mechanistic knowledge; at this point our performance is inferior to their devices, so the elites easily persuade us that the only way to improve is by using devices; as we get to depend on devices, the only knowledge we acquire is the mechanistic knowledge required to operate them; so our skills remain below the level of devices, and we believe that we must continue to depend on them. The only way to escape from this trap is by developing non-mechanistic knowledge, and thus becoming superior to the devices. But this takes time, and time is precisely what we do not have if we squander it on a preoccupation with devices. As long as we trust the elites, therefore, we are condemned to using only the *mechanistic* capabilities of our minds; we are condemned, in other words, to staying at novice levels forever.

Let me put this differently. To control our life, the software elites must induce a state of permanent ignorance and dependence. And they achieve this by persuading us to trust their mechanistic expedients – concepts, theories, methods, devices – while these expedients can rarely solve our real problems. Consuming our time by keeping us preoccupied with their expedients is a critical factor in the process of domination, because the elites must prevent us from using this time to develop our minds. And promoting worthless expedients is an integral part of this process: they wouldn't give us useful ones even if they could. Only expedients that do *not* work can be employed to consume our time; only by *not* solving our problems can they add to our spurious, software-related preoccupations. No domination would be possible if we were asked to depend on the elites only in those few situations where their expedients are indeed superior to our minds (that is, where a complex phenomenon can be usefully approximated with simple structures).

Promoting mechanistic expedients, thus, ensures our continued dependence in two ways at once: by restricting our knowledge and skills to levels even lower than those attained by the expedients, and by consuming our time with the endless preoccupations generated by the expedients.

2

It may be useful to recall our software preoccupations, although we are already spending so much time with them that they are well known. Installing new software, seeking technical support, downloading updates, studying lists of

“frequently asked questions,” checking the latest notes on a website, reading computer magazines, discovering “undocumented features,” running virus protection utilities, printing online manuals, trying to get different pieces of software to work together – these are some of the activities we must perform when involved with software. But these are only the incidentals. We must also include the time required to learn to use the software (the features and options we have to assimilate, how to specify and combine them, keeping up with changes from one version to the next), and the time we take to actually use it, once we get to depend on it.

These activities require almost exclusively mechanistic knowledge: they consist of isolated and fairly simple tasks, which cannot help us to develop an important body of knowledge or skills. We note this in that almost everyone, regardless of age or experience, has to deal with the same kind of problems; and almost everyone manages to solve them. We also note it in that, no matter how many of these problems we faced in the past, we will still face similar ones in the future. In other words, the proportion of time we must devote to software-related problems does not decrease significantly with experience.

Anyone who encountered software-related problems is familiar with the feeling of satisfaction experienced when finally uncovering the answer. The answer is usually a simple fact; for instance, learning that we must select one option rather than another in a particular situation. But instead of being outraged that we had to spend time with an activity so simple that we could have performed it as children, we perceive it as an essential aspect of our work, so we believe that we have learned something important. Although we don't think of this activity as a form of amusement, we experience the satisfaction of solving a puzzle. And even if it is true that we must now spend a great part of our time solving puzzles instead of addressing real problems, it is significant that these are *trivial* puzzles, demanding only a fraction of our mental capabilities. Clearly, there is no limit to the number of software-related puzzles that we can find, and hence the time we must take to deal with them, if we agree to depend on concepts and products that cannot solve our real problems to begin with.



Any activity, method, or tool entails some incidental preoccupations, so we cannot expect to benefit from software without investing some time; and we may even have to spend part of this time dealing with trivial issues. Thus, what I am trying to show here is *not* that our collective preoccupation with software is too great relative to the benefits we derive from it. Such deficiency we could attribute to the novelty of software and to our inexperience. We could then

conclude that this condition is transient, and that we will eventually become as able in our software pursuits as human beings can be – just as we have become in other domains.

What I am trying to show, rather, is that this interpretation is wrong, that our incompetence is getting worse not better, that our software preoccupations do not reflect a natural process of intellectual evolution. On the contrary: the incompetence is deliberately fostered by the software elites as part of a monstrous plan of domination founded on our mechanistic delusions, and made possible by our growing dependence on computers – a plan whose goal is to degrade the mind of every human being on earth.

Our continued ignorance in software-related matters – programming, in particular – is essential in this plan of domination, because software knowledge, like linguistic knowledge, is related to all other types of knowledge. *Software* ignorance and dependence, thus, are only the means to bring about *total* ignorance and dependence. It is in order to induce this collective ignorance and dependence that the software elites are exploiting our mechanistic delusions, and the consequent software delusions. For, as long as we believe that knowledge and skills can be replaced with mechanistic concepts, we will inevitably conclude that we must depend on organizations that produce devices based on these concepts – just as we depend on organizations that produce appliances, detergents, or electricity.

In reality, to succeed in software-related activities – programming, in particular – we need *skills*. And, like other skills, these new skills depend on our own capabilities and experience. Also like other skills, they demand the full capacity of the mind, and, to attain expertise, many years of learning and practice.

The software elites are promoting the notion of knowledge substitutes precisely because these substitutes are worthless. It is precisely because they cannot replace skills, and hence fail to solve our problems, that we constantly need new ones and spend so much time with them. By consuming our time with the petty preoccupations generated by these substitutes, the elites are preventing us from developing skills, thus ensuring our continued incompetence and dependence.

Were we permitted, as society, to develop software skills as we develop skills in other domains – were we permitted, in other words, to attain the highest level that human minds can attain in software-related matters – the issues of incompetence and waste of time would not even arise. We would then be, quite simply, as good in these new skills as we can possibly be; and we would take as much time with our software preoccupations as is justifiable. This is how we progressed in other domains, and there is no reason to view software and programming differently. It is unlikely that we have already reached the highest

level, or that we are advancing in that direction, since we are using now mostly *mechanistic* knowledge; in other domains, it is our *non-mechanistic* capabilities that we use when we attain expertise. The software elites can persuade us to prefer their mechanistic substitutes to our own minds only because, as society, we have no idea how good we can actually be in software-related matters: we never had the opportunity to find out.

All skills – whether easy or difficult, acquired early in life or later – entail the same mental processes. Interpreting visual sensations, recognizing social contexts, diagnosing diseases, playing musical instruments, flying airplanes, repairing appliances, teaching children, managing warehouses – we can acquire almost any skill, but the only way to acquire it is by performing the activities involved in that skill, and by allowing our mind to discover the complex knowledge structures which constitute that skill. For no other skills can we find an elite that prevents us from using the full capacity of our mind, or forces us to use methods and devices instead of expanding our knowledge, or redefines expertise to mean expertise in the use of substitutes for expertise. From all the skills we can acquire, only those associated with software and programming seem to have engendered such notions, and the reason is simple: these skills are *more complex* than the others, and hence misunderstood. They are so complex, in fact, that they permit us to use our mind and view the world in entirely new ways. They are comparable in scope only to our linguistic skills.

So it is the complexity itself that allows charlatans to deceive us. We became infatuated with software too quickly, without taking the time to appreciate the true meaning and the implications of software knowledge. We still do not understand what can happen when a society depends on software while software is controlled by an authority. This is why we continue to accept the absurd notions promoted by the software elites; in particular, the notion that we must depend, in all software-related affairs, on systems produced by software companies. It is precisely because software knowledge is so difficult that we are tempted by theories which tell us that we can enjoy the benefits of software without taking the time to develop software knowledge.

3

Our preoccupation with *ease of use* deserves a brief analysis. Software – applications, development systems, utilities – is always promoted with the claim that it is easy to use. I want to show, however, that the belief in easy-to-use software is a mechanistic delusion. The notion “easy to use” is, strictly speaking, meaningless.

Like any tool or device, a piece of software cannot be any easier to use than whatever effort is required to accomplish a given task. The only sensible claim, therefore, is that it is *well-designed*. A lathe, for example, even if well-designed, is necessarily more difficult to use than a chisel. And so it is with software: all we can expect of a particular business application, or a particular development tool, is that it be well-designed. Once this requirement is fulfilled, the notion “easy to use” becomes irrelevant: that software will be as easy or as difficult to use as software can be in a particular situation.

Now, we see the claim “easy to use” for *all* types of software – for business and for home, for programmers and for end users. We never see software described, for example, with the warning that we need much knowledge, or many months of study and practice, in order to enjoy its features. Thus, as we are becoming dependent on software in practically everything we do, if all this software is also easy to use, we reach the absurd conclusion that human beings will never again have to face a challenging situation.

The delusion of easy-to-use software becomes clearer if we recall the other quality commonly claimed for software – *power*. Just as all software devices are said to be easy to use, they are also said to be powerful. The two qualities are often claimed together, in fact, as in the famous phrase “powerful yet easy to use.” By combining the two qualities, the following interpretation presents itself: we believe that software devices embody a certain power, and we perceive ease of use as the ease of invoking this power.

The only power that can inhere in a software device is its built-in operations; that is, higher levels of abstraction for our starting elements. And it is the higher levels that also make the device easy to use. The power and ease of use are illusory, however: high starting levels make the device convenient when our needs match the built-in operations, but awkward or useless otherwise.

We saw this with the language analogy: we have *less* power when starting with ready-made sentences; we must start with *words* if what we want is the capability to express any conceivable idea. Similarly, if true software power is the capability of a system to handle any situation, the only way to have this power is by starting with *low-level* entities. Like claiming ease of use, therefore, claiming power for a software device is nonsensical: what is usually described as power is the exact opposite of software power. Moreover, since ease of use can be attained only by providing higher starting levels, and hence by *reducing* the power of the device, claiming both power and ease of use at the same time is especially silly.

Putting all this together, it is obvious that the software elites want us to think of software as an assortment of devices that have the power to solve our problems, while all *we* have to do is *use* them. The only thing left for us to do from now on is operate software devices; and this we can learn in a matter of

hours. This notion is absurd, as we just saw, but we enthusiastically accept it. The elites are plainly telling us that we will no longer have the opportunity to use the full capacity of our minds, that our sole responsibility will be to perform tasks so simple that anyone can learn to perform them in a short time. But instead of being outraged, we welcome this demotion; and, to rationalize it, we interpret our diminished responsibility as a new kind of expertise.



The elites also claim that software devices will enhance our creativity, by taking over the dull, routine activities. With higher starting levels, the elites tell us, we can reach the top element of a given structure much sooner. Why waste our time and talents with the details of the low levels, when the functions built into these devices already include all the combinations of low-level elements that we are likely to need? When starting from low levels we squander our superior mental capabilities on trivial and repetitive tasks; let the computer perform this tedious work for us, so that we have more time for those tasks demanding creativity. Just as successful managers and generals deal only with the important decisions and leave the details to their subordinates, we should restrict ourselves to high-level software entities and leave the details to the computer.

It is easy to show the absurdity of these claims. We are told to give up the details, and to use instead ready-made entities, so that we have more time for the important work, more time to be creative. But our work *is* the development of high-level entities from low-level ones. In one occupation after another, the software elites are redefining the concept of work to mean the act of combining the high-level entities provided by their devices. To be creative, however, we must be able to arrive at *any one* of the possible alternatives; and this we can do only by starting with *low-level* entities. Moreover, we are offered software devices in *all* fields of knowledge, so we cannot even hope that the time we perhaps save in one type of work will permit us to be more creative in another.

Returning to the language analogy, if a writer used ready-made sentences instead of creating new ones, starting with words, we would study his work and recognize that he is not being *more* but *less* creative. Clearly, fewer ideas can be expressed by selecting and combining ready-made sentences than by creating our own, starting with words. And this is true for all types of knowledge: the higher the level we start with, the greater the effect of reification and abstraction, and the fewer the alternatives for the top element. So it is absurd to claim that we can be more creative by avoiding the low levels, seeing that it is precisely the low levels that make creativity possible.

Managers and generals who make good decisions only *appear* to start from

high levels. In reality, their decisions involve knowledge structures that interact at low levels, at the level of details. But this is largely intuitive knowledge, so all we can observe is the top element of the complex structure; that is, the final decision. (See “Tacit Knowledge” in chapter 2.) They developed their knowledge over many years, by dealing with all structures and all levels, low and high. This is the essence of personal experience. Were their knowledge limited to the high levels, to those selections and combinations that can be observed, then anyone could quickly become a successful manager or general – simply by learning to select and combine some high-level concepts.

This delusion is also the basis of the software devices known as expert systems – one of the sillier ideas in artificial intelligence. Expert systems claim that it is possible to capture, in a specially structured database, the knowledge possessed by a human expert in a given domain. The database consists of answers that the expert provides to certain questions – questions formulated so as to simulate various decision-making situations. Then, for a real problem, simply by interrogating the system, anyone should be able to make the same decisions that the expert would make. The fact that such devices are being considered at all demonstrates the degradation in the notions of expertise and responsibility that we have already suffered. As we saw in “Tacit Knowledge,” expertise is the level where a person does *not* have to rely on rules, methods, and databases of facts (see pp. 157–158). Thus, the device can capture only the *mechanistic* aspects of the expert’s knowledge; and consequently, a person using it will not emulate an expert but a novice.

Another claim we see is that software devices enhance our creativity by giving us new forms of expression. And this claim, too, is empty. Software does indeed allow us to express ourselves and to view the world in new ways, as does language. But, as in the case of language, we can only enjoy this quality if we develop our structures starting with *low-level* entities. For, only then can we discover all possible interactions between the software structures, between software structures and the other structures that exist in the world, and between software structures and the knowledge structures present in our minds. If we get to depend on software devices, and hence on *high-level* software entities, we will not only fail to develop all possible alternatives in the new, software-related matters, but we will lose alternatives in the knowledge and skills that we had in the past. In the end, we will have *fewer* alternatives than before, *fewer* ways to express ourselves. Thus, far from *enhancing* our creativity, software devices are in fact degrading our minds, by forcing us to spend more and more time with activities requiring largely mechanistic knowledge.



Power and ease of use, thus, are specious qualities. The elites tell us that software devices can have these qualities because they want to replace our traditional conception of expertise with a dependence on these devices. They want us to believe that all the knowledge that matters inheres now in software devices, so all *we* have to know is how to operate them. The implicit promise is that, thanks to these devices, we don't need to know anything that we don't already know – or, at least, anything that we cannot learn in a short time.

So the elites are downgrading our conception of expertise by reducing to a minimum the range from novice to expert. If all we have to know is how to operate software devices, the difference between novice and expert is just the time taken to acquire this knowledge. Where we thought that one needs many years of study and practice to attain expertise in a difficult field, we are told that this type of knowledge is obsolete. The propaganda makes it seem modern, sophisticated, and glamorous to perform a task by operating a software device, and unprofessional or old-fashioned to perform it by using our minds. Consequently, we are according more importance to our methods and tools than we do to the results of our work. Increasingly, we are judging a person's knowledge and skills by his acquaintance with software devices, instead of his actual capabilities and accomplishments. Increasingly, it doesn't even matter what the results are, as the main criterion for assessing a professional activity is whether the person is using the latest software devices.

4

Recall the pseudosciences we studied in chapter 3. I stated there that our software delusions belong to the same tradition, that they are a consequence of the same mechanistic culture. With our software theories we are committing the same fallacies as the scientists who pursue mechanistic theories in psychology, sociology, or linguistics. When we waste our time with the spurious problems generated by our mechanistic software concepts, we are like the scientists who waste their time studying the mechanistic phenomena created by reifying human phenomena. Just as those scientists cannot explain the complex phenomena of mind and society by explaining separately the simple, mechanistic phenomena, *we* cannot solve our complex social or business problems by solving the simple, software-related problems.

These mechanistic delusions I have called *the new pseudosciences*, and we saw that they are similar to the traditional pseudosciences – astrology, alchemy, and the rest. They are similar in that they too are systems of belief masquerading as scientific theories, and they too are based on hypotheses that are taken as unquestionable truth. In the case of the new pseudosciences, the

hypothesis is that mechanism can provide useful explanations for complex phenomena – for phenomena involving human minds and societies, in particular. The mechanists are, in effect, today’s astrologers and alchemists: respected thinkers who attract many followers, even though their theories do not work.

Before we had software, it was only in the academic world that one could spend years and decades pursuing a mechanistic fantasy. One could hardly afford to fall prey to mechanistic delusions in business, for instance. But through software, the ignorance and corruption engendered by mechanistic thinking is increasingly affecting the entire society: corporations, governments, individuals. Through software, we are all asked now to accept fantastic mechanistic theories – theories that promise to solve our problems with practically no effort on our part. Through software, the entire society is returning to the irrationality of the Dark Ages: we are increasingly guided by dogmas instead of logic, by beliefs instead of reason.

When we believe that a software device can replace knowledge, skills, and experience, we are committing the same mistake as the scientists who believe that mechanistic theories can explain human intelligence and social phenomena. So if all of us now, not just the academics, are wasting our time with pseudoscientific theories, we must ask ourselves: Can we afford this corruption? Can our civilization survive if *all* of us engage in futile mechanistic pursuits? When mechanistic theories fail in the academic world, the harm is limited to a waste of resources, and perhaps a lost opportunity to improve our knowledge through better theories. But what price will we pay if we create a society where *all* theories fail?

As we are modifying our values and expectations to fit the mechanistic software ideology, we are adopting, in effect, mechanistic theories – theories on our capabilities as human beings, or on our responsibilities as professionals. And since these software-based theories suffer from the same fallacies as the traditional mechanistic theories, they too must fail. But what does it mean for *these* theories to fail? Since what they claim is that we can accomplish more by depending on software devices than by developing our minds, a failure of *these* theories means that we are making a wrong decision about ourselves: we mistakenly assume that our minds can be no better than some mechanistic expedients. Thus, when we decide to leave our non-mechanistic capabilities undeveloped and to depend instead on mechanistic expedients, we are causing, quite literally, a reversal in our intellectual evolution: we are choosing to degrade our conception of intelligence to a mechanistic level, and to create a world where there is no need or opportunity to exceed this level.

Let us interpret the new pseudosciences in another way. The equivalent of a world where we depend on software while being restricted to mechanistic

software theories is an imaginary world where the *traditional* mechanistic theories – those explaining minds and societies – actually work. Since these theories fail to explain our *real* intelligence and behaviour, in the imaginary world we would have to alter minds and societies to fit the theories. To comply with the linguistic theory of universal grammar, for example, we would restrict our sentences, and the associated thoughts, to what can be depicted with exact diagrams and formulas; similarly, to comply with behaviourism or cognitive science, we would restrict our behaviour and mental acts to patterns that can be precisely explained and predicted; and to comply with the theories of structuralism or the social sciences, we would restrict our institutions, customs, and cultures to activities that can be described mathematically.

These theories reflect the diminished view that mechanists have of human beings – the view that our acts can be explained with precision, because our capabilities are like those of complicated machines. The scientists who invent and promote mechanistic theories wish them to work, of course. But the theories can work only if we are indeed like machines, so we must conclude that these scientists *want* us to be like machines. And if we, the subjects of these theories, also wanted them to work – if we agreed, as it were, to satisfy the wish of their authors – we would have to restrict our capabilities to what these theories can explain and predict. In other words, we would have to mutate into automatons.

What has saved us from this fate so far is not wisdom – for, if we had that wisdom we would have abandoned the mechanistic ideology already – but the fact that none of these scientists had the power to make us conform to their theories. Through software, however, it has finally become possible for the mechanists to realize their dream: a world where human beings can be designed and controlled as successfully as we design and control machines. The world that we can only *imagine* through the traditional mechanistic theories, we are actually creating through our mechanistic *software* theories. Whereas we can still think, learn, speak, and behave while ignoring the mechanistic theories of mind and society, we are forced to create and use software according to mechanistic theories. But if we are to depend on software in all aspects of our life – including those aspects studied by the theories of mind and society – then by following mechanistic software theories we *are*, in effect, mutating into the automatons that the mechanists wish us to be.

Remember, though, that it is not software dependence in itself that is harmful. On the contrary, if we were permitted to use it freely, as we use language, software would *enhance* our mental capabilities, as does language. The danger lies in the dependence on software while software knowledge is restricted to its mechanistic aspects – a policy intended to prevent us from using the full capacity of our minds.



The decision we are making now is more than a choice; it is a commitment. As individuals and as society, we are making a commitment; namely, to invest in software expedients rather than our minds.

As individuals, we reaffirm this commitment when we consent to depend on software devices that are inferior to our own minds; when we spend time solving a specious, software-related problem, instead of expanding our knowledge to deal with the real problem; and when we degrade our conception of professionalism and responsibility, from the utmost that human beings can do, to merely knowing how to use software devices. As society, we reaffirm this commitment when our corporations and governments, instead of encouraging their workers to develop expertise, spend vast amounts of money on projects that increase their dependence on the software elites.

As individuals, if we are wrong, our knowledge in, say, ten years will not be much greater than what it is at present. We will waste that time acquiring worthless bits of knowledge; specifically, knowledge of ways to avoid the need for real knowledge. If we make this choice, of course, we will be unable to recognize our own ignorance in ten years; so for the following ten years we will make the same choice, and so on, and we will remain for the rest of our lives at the present level. As society, if we are wrong, within a few decades we will be where we were centuries ago: in a new dark age, ruled by elites that know how to exploit our ignorance and irrationality.

The decision we are making now is a commitment because we cannot choose both alternatives. If software mechanism is our decision, we will need only mechanistic capabilities; so we will leave our superior, non-mechanistic capabilities undeveloped. If we are wrong, we cannot reverse this decision later: if we choose the mechanistic alternative, in any domain, we will not practise; and practising is the only way to develop non-mechanistic knowledge. If we lose our appreciation of non-mechanistic knowledge, we will forget, in one occupation after another, that we *are* capable of more than just following methods and operating software devices.

This is precisely what has happened in the domain of programming. The superior alternative – personal knowledge and skills – is always available, in principle: any programmer, any manager, any company, could choose to ignore the official software ideology and treat programming as we do the other professions. Yet, despite the evidence that programming aids and substitutes are inferior to human expertise, we continue to trust the software elites and their mechanistic theories. In the domain of programming, we have already lost our appreciation of non-mechanistic knowledge.

Software Charlatanism

Software exploitation, we saw, plays on our mechanistic delusions; namely, on the belief that problems requiring complex knowledge can be broken down into simpler problems, which can then be solved mechanistically. When the software charlatans tempt us with the promise of easy answers to difficult problems – answers in the form of software devices – what they do is tempt us to commit the mechanistic fallacies, reification and abstraction. For, only if we commit these fallacies will we believe that software devices can be a substitute for the complex, non-mechanistic knowledge required to solve those problems.

In the present section, we will examine how the mechanistic delusions manifest themselves in various software-related activities; that is, how the two mechanistic fallacies lead to *software* delusions, and how these delusions are being exploited by the software elites.

The Delusion of High Levels

1

We are deceived by the mechanistic software theories because we like their promise. The promise, essentially, is that methods and devices simple enough to be used by almost anyone can be a substitute for programming expertise. And we believe this promise because we fail to see that to accept it means to commit the two mechanistic fallacies, reification and abstraction.

The software theories appear to make programming easier because they treat applications, or the activities involved in creating applications, as separable into independent aspects. Each aspect of an application is one of its processes, one of the structures that make it up. Thus, each subroutine together with its uses, each database field or memory variable together with the associated operations, each business rule or programming method, can be seen as a simple structure. But these structures are not independent. Although in our imagination we can separate them, in reality they share their elements (the software entities that make up the application), so they interact. (See “Software Structures” in chapter 4.)

The software theories invite us to reify programming and applications because, once we have independent structures, they can tempt us to start from higher levels of abstraction within each structure. By the time we commit both fallacies, the concept of programming, and the resulting software, are

impoverished: many of the functions we could implement before are no longer possible.

All software theories, then, make the same claim: the task of programming can be simplified by starting the development process from higher-level software elements; and we can accomplish this by allowing various expedients – methodologies, software tools, built-in operations – to act as substitutes for the knowledge, experience, and work necessary for creating the lower levels. But this would be possible only if applications consisted of independent structures.



To demonstrate the fallacy of high starting levels, let us analyze a specific situation. A common requirement, found in most business applications, is to access individual fields in database files. The application's user may need to see the phone number or the outstanding balance of a customer, or the quantity in stock of a certain part; or he may need to modify the address of a customer, or the description of a part. These operations may involve several files; for example, a customer is displayed together with its outstanding invoices, or a part together with its sales history. Typically, the user specifies some values to identify the records: customer number, invoice number, range of dates, etc. The program displays certain fields from those records, and the user may modify some of them. These can be isolated fields, but most often they are groups of fields logically associated with specific functions: inventory control, financial information, shipping activity, etc. If we also include such options as adding new records and deleting existing ones, we may refer to this category of operations as *file maintenance* operations.

Now, file maintenance operations constitute fairly simple programming tasks. Moreover, much of this programming is very similar in all applications. So it is tempting to conclude that we can replace the programming of file maintenance operations with a number of high-level software elements – some built-in procedures, for example. We should then be able to generate any file maintenance operation by combining these high-level elements, rather than starting with the individual statements and operations of a traditional programming language. I want to show, though, that despite the simplicity and repetitiveness of file maintenance programming, it is impossible to start from higher-level elements.

The illusion of high levels arises when we perceive software as a combination of separable structures, or aspects. There are at least two aspects to the file maintenance operations: database operations and user interface operations. So, to keep the discussion simple, let us assume that these two aspects are the only important ones. If we think of each aspect separately, it is quite easy to

imagine the higher levels within each structure, and to conclude that we can start from higher levels. We may decide, for example, that most database operations can be generated by starting with some built-in procedures that let us access specific records and fields; and most interface operations, with some built-in procedures that display individual fields and accept new values. Thus, by specifying a few parameters (file and field names, index keys, display coordinates, etc.), we should be able to generate, simply by invoking these procedures, most combinations of database operations, and most combinations of interface operations.

We commit the fallacy of abstraction, however, if we believe that the alternatives possible when starting from higher levels are about the same as those we had before. The temptation of high levels is so great that we are liable to perceive our needs as simpler than they actually are, just so that we can rationalize the reduced flexibility. It takes much experience to anticipate the consequences of the restriction to high levels. For, it is only later, when the application proves to be inadequate, when important requirements cannot be met and even simple details are difficult to implement, that the impoverishment caused by abstraction becomes evident.

But abstraction became possible only through reification – only after separating the two structures, database and user interface. And reification causes its own kind of impoverishment. We can indeed view file maintenance operations from the perspective of either the database or the interface operations, but only in our imagination. In reality, the file maintenance operations consist of *both* the database and the interface operations. Separately, these operations can indeed be represented as simple structures, because we can identify most of their elements and relations. But when part of an application, these operations interact, giving rise to a complex structure. It is this complex structure that constitutes the real file maintenance operations, not the two imaginary, reified simple structures.

Reification impoverishes the complex structure by destroying the interactions between its constituent structures. When we lose the interactions, we also lose many alternatives for the top element of the complex structure. Each alternative at this level represents a particular file maintenance operation, which may be required by some application. Certain interactions are still possible, of course – those that can be generated by combining the high-level, built-in procedures. The alternatives resulting from these interactions are the file maintenance operations that can be implemented even after reification and abstraction. Most interactions, however, take place at low levels, so they can only be implemented with such means as statements, variables, and conditions; that is, with programming languages. And these interactions are no longer possible once we lose the lower levels. The two fallacies, thus,

contribute together to the impoverishment of the complex structure that is the application. They are usually committed together, and it is seldom possible, or necessary, to analyze them separately.

To appreciate why it is impossible to eliminate the low levels, all we have to do is think of the details that a programmer faces when implementing a typical file maintenance operation. Thus, the user may want to see only some of the fields at first, and then various other fields depending on the values present in the previous ones; or he may need to scan records, forward or backward, rather than ask for specific ones; in one application the user may want to see detailed information, in another only a summary; in one situation some of the fields may always be modified, in another the fields may be modified only under certain conditions; in some applications, modifying a field must produce a change in other fields, and perhaps in other files too; and so on.

Clearly, the number of possible requirements, even for relatively simple operations like file maintenance, is practically infinite. But the important point is that this variety, and the details that make up these requirements, entail the low levels of *both* the database and the interface operations. To implement a particular requirement, therefore, we need not only low-level software elements in both kinds of operations, but elements that can be *shared* by these operations; in other words, exactly what abstraction and reification would *prevent* us from creating. For example, to display a field depending on the value of another field, we must formulate conditional statements involving particular fields and display operations; and to display details from one file along with the summary of another, we must create a small piece of software that reads records, accesses fields, performs calculations and comparisons, and displays values.

Each requirement reflects a particular file maintenance operation; each one is, therefore, an alternative value for the top element of the complex structure formed by the interaction of the database and display operations. If we agree that a programmer must be able to implement *any* file maintenance operation, and hence to generate *any* alternative, it is obvious that he must be able to create and combine all the low-level elements forming the database operations, and all the low-level elements forming the display operations (and probably other low-level elements and operations too). The use of low levels helps us avoid both fallacies: it lets us generate all the alternatives within each structure, *and* the alternatives resulting from the interaction of the two structures.

Since each alternative is unique, no matter how many alternatives we have already implemented, or are available through built-in procedures, the next application may still have to be programmed starting with low-level elements. Only naive and inexperienced practitioners believe that they can have the versatility of the low levels while being involved only with high levels. In

reality, the simplicity promised for high-level operations is achieved precisely by reducing the number of alternatives. It takes the experience of many applications to recognize, in a given situation, whether we can or cannot give up the low levels.

2

File maintenance was only an example, of course. A major application comprises *thousands* of aspects, most of them more involved than a database or display operation. Besides, we seldom encounter situations where only two aspects interact, as in our simplified file maintenance example. Even there, to discuss realistic situations we had to consider, in addition to the database and display operations, various business practices. These practices are themselves aspects of the application, so they add to the number of structures that must interact. We saw this, for instance, when I mentioned the small piece of software that accesses records and fields, performs calculations and comparisons, and displays values: a small element comprising just a few statements must be shared, nevertheless, by several processes – database, display, and one or more business practices – because this is the only way to implement an operation that involves these processes.

It is hardly necessary, therefore, to demonstrate the need for low levels in real applications, in situations involving thousands of aspects, after showing the need for them even in situations with two aspects. Rather, what I want to show is how the mechanistic fallacies, and the software delusions they engender, lead to software charlatanism. All forms of software exploitation are based, ultimately, on the delusion of high levels that we have just examined.



The deception starts when we are offered some software that promises to enhance our capabilities; namely, software that will allow us to accomplish tasks requiring knowledge that we, in fact, lack. We are promised, in other words, that simply by operating a software device we will be in the same position as those whose skills and experience exceed ours. The promise, it must be emphasized, is not that we will quickly acquire the missing knowledge. On the contrary, the promise is specifically that we don't need to learn anything new: the power to perform those tasks resides in the device itself, so all we need to know is how to operate it.

As programmers, we are offered various tools, development environments, and database systems. We are told that these devices will enable us to create,

quickly and easily, applications which otherwise would take us a long time to program, or which are too difficult for us to program at all. The promise, therefore, is that these devices will function as substitutes for programming expertise: through them, we will achieve the same results as programmers who have been developing and maintaining applications successfully for many years.

As users, we are offered various productivity systems, or office systems. We are told that these devices will solve our business problems directly, eliminating the need for programming. Or, we are offered ready-made applications or pieces of applications, and we are told that they will enable us to manage our business just as we would with custom applications created specially for us.

For programmers as for users, the promises are supported with the explanation that the software devices offer higher levels of abstraction: they simplify development by allowing us to start from higher-level elements, bypassing the difficult and time-consuming task of creating the lower levels. The higher-level elements appear in a variety of forms, but, essentially, they are built-in operations or ready-made pieces of software.

No matter what form the higher levels take, the underlying assumption is the same: the work involved in creating a software application is similar to a manufacturing project, so the application can be seen as a neat structure of things within things – parts, modules, subassemblies. Thus, as in manufacturing, the larger the building blocks, the faster we will complete the project. We should avoid programming, therefore, and start instead with the largest modules and subassemblies available: software entities that already contain the lower-level parts. The use of large building blocks benefits us in two ways: by speeding up the software manufacturing process, and by demanding lower skills. Less time, less knowledge, and less experience are needed to assemble a software structure from modules, than to design and build it from basic components. In the extreme case of ready-made applications, the manufacturing process is eliminated altogether: the starting level is then the top element itself, the complete application.

If software exploitation begins with the lure of high levels, the next stage is, needless to say, the disappointment. As programmers, we still cannot create complex and reliable applications; as users, we still cannot manage our affairs as we hoped. The software devices do provide the promised higher levels, and we can perhaps accomplish some tasks that we could not have accomplished without them. What we find, rather, is that the higher levels are rarely beneficial. If we want to start from higher levels, we must give up the flexibility afforded by the low levels. If we want the benefits of built-in operations and ready-made modules, of less work and easier challenges, we must be content with a fraction of the alternatives otherwise possible.

Unfortunately, only rarely is this practical: only rarely can we restrict our affairs to the few alternatives provided by the software devices. The greater the promised benefits, the higher must be the starting levels, and the more severe the reduction in alternatives. The deception, thus, consists in promoting the benefits of higher starting levels while masking their concomitant drawbacks.

We adopt these devices and become dependent on them because we are seduced by slogans like “powerful” and “easy to use.” We fail to see that these two qualities are contradictory: easy-to-use devices can be powerful only if we redefine power to mean, not the ability to implement *any* operations, but the ability to implement *some* operations easily. If ease of use is claimed, then, clearly, we are expected to do little; so all the power must inhere in the devices themselves, in their built-in capabilities. This means that they may perform well those operations that are built in, but they cannot perform other operations at all; and no device can have all conceivable operations built in.

The only way to implement *any* operations that may be required is by starting with low-level elements. So the software charlatans must provide the low levels if we are to use their devices at all. Their challenge, therefore, is how to reinstate the low levels, and how to make us start from these low levels, while we continue to believe that we are working at high levels. And they do it by implementing the low-level features *within* the high-level environment, as *extensions* to the high-level operations.

The low levels were always available to us – in the form of traditional programming languages, for example. So, if low levels are what we need, there is nothing the elites can give us that we did not have all along. The theories and methodologies, the programming tools and fourth-generation languages, the database and reporting systems, serve in reality the same purpose: they provide some of the low-level elements we need, and the means to link software structures, while pretending to be high-level environments.

The third stage in the process of exploitation, then, is the reinstatement of the low levels. To make their devices useful, the elites must restore the very concept that the devices were meant to supersede. Any device that does *not* provide this functionality is eventually abandoned and forgotten, even by the naive people who believed the original claims, simply because it is useless. (The full-fledged CASE environments, which actually tried to materialize the fantasy of creating entire applications “without writing a single line of code,” are an example.)

We will waste no time, thus, examining the devices that do *not* restore the low levels. Let us treat them simply as fraudulent products, no different from the other forms of deception employed by charlatans to exploit gullible people – weight-loss contraptions, back-pain remedies, money-making schemes, and the like. As explained earlier, it is not the traditional means of exploitation, but

the new form of domination, that concerns us: the use of software to consume our time and prevent us from gaining knowledge and experience. Eliminating the low levels and then restoring them in a different and more complicated form is an important factor in this domination, as it permits the elites to destroy software freedom and to establish the dependence on their devices. And we are fooled by these charlatans because the devices are based on software theories invented in universities, and described as “scientific.”

Recall the principles of demarcation between science and pseudoscience, which we studied in chapter 3. Mechanistic software theories claim that we can create applications by starting with high-level elements. So, when this idea proves to be worthless and the charlatans “enhance” their devices by restoring the low-level capabilities, what they do in reality is turn the falsifications of those theories into new features. And this, we saw, is the stratagem through which fallacious theories are rescued from refutation. Thus, mechanistic software theories are intrinsically pseudoscientific.

Here are examples of software devices that were enhanced by restoring the low levels: The so-called fourth-generation languages started by promising us a higher level than the traditional, third-generation languages; but the only way to make them practical was by restoring the features found in the traditional languages (loops, conditions, individual variables, etc.). The relational database systems started by claiming that the database can be treated as separate structures, interacting only at high levels with the other structures of the application; but they became practical only after adding countless new features, and whole programming languages, in order to restore the low-level links between these structures (for instance, the capability to access individual records directly, through file scanning loops). Systems like report writers and spreadsheets started by claiming that their high-level features are adequate for our requirements; but they ended up incorporating many traditional features, and even programming languages, in order to provide the low-level operations needed in real-world situations.

To summarize, high-level environments that restore the low levels exploit us in two ways. First, we get to depend on some new and complicated programming methods, arising from the idea of using low-level elements as an extension to high-level ones. The traditional method – creating high-level elements from low-level ones – is simple and natural; it follows a concept we all understand intuitively, and confers complete programming freedom. The new methods, on the other hand, are contrived – absurd and unnecessary; their purpose is to maintain the illusion of high levels, and to induce dependence on proprietary development systems. Second, these systems provide only *a few* of the low-level features available through the traditional methods, only the minimum necessary to fool us. So they remain, essentially, high-level

environments, lacking the versatility of a general-purpose language. Each low-level feature is presented as a powerful enhancement, and this obscures the fact that these features are merely a more complicated version of features we always had – in the traditional programming languages.

3

The charlatans promise us power, but all they can give us is higher levels of abstraction. So, once they persuade us to adopt their devices, they must restore the essential low levels, while continuing to promote the devices as high-level environments. It is precisely because they are designed to deceive us – to prevent us from noticing that we are working at *low* levels – that the devices end up so complicated and inefficient.

The so-called non-procedural languages, for instance, are promoted with the claim that we only need to tell the computer now *what* to do, not *how* to do it. This sounds like a novel programming concept, as if we could almost talk to the computer and tell it what we need done. In reality, these languages merely incorporate a number of high-level elements in the form of built-in operations. And this concept is available in any programming language in the form of libraries of subroutines: ready-made functions providing levels of abstraction that are as high as we want.

But, whether we use subroutines or non-procedural languages, some of our starting elements must still be at low levels, because it is impossible to implement all conceivable requirements by relying entirely on ready-made, high-level elements. The non-procedural languages look impressive when all we need is one of their built-in operations (these are the examples we see in textbooks and in advertisements, of course), but are more awkward than a traditional language in any other situation. This is true because, in order to make them appear as novel concepts, their authors must make them different from the traditional programming languages; and this also makes them more complicated.

It is common in these languages, for example, when what we need is not just one of the built-in operations, to find long and tangled statements. A procedure that in a traditional language involves conditions, loops, and the use of variables, may turn up in a non-procedural language, in an extreme case, as *one* statement. But, while being perhaps shorter than the procedure, the statement is not, in fact, a higher-level entity: since its clauses must be specified with precision and in detail, they do not constitute a higher level of abstraction.

Thus, in SQL (the most popular database language), we often see statements containing more than a dozen lines, when a number of related database

operations must be specified together. These long statements can become extremely complicated, as the programmer is forced to cram into one expression a whole series of related operations. Instead of the familiar structure of loops and conditions found in traditional languages, and which an experienced programmer understands intuitively, we have now an artificial and unwieldy set of specifications. But because the definitions, loops, and conditions are no longer manifest, this complicated piece of software is unlike a traditional language, so we can delude ourselves that what we are doing is no longer programming: “we are only telling the computer what to do, not how to do it.”

What we are telling the computer is, however, the same as before. First, the level of abstraction is about the same as in a traditional language (this, after all, is why we needed SQL, why we could not simply use the high-level relational operations). Second, the resulting statements are still the reflection of many structures, which interact and must be kept in the mind simultaneously. In other words, all the difficulties we had before are still there; and because we wanted to *avoid* programming, we must now cope with these difficulties through programming means that are more complicated and less efficient than the traditional ones. (We will examine the SQL fraud in “The Relational Database Model” in chapter 7; see pp. 794–801.)



Let us look at another concept that promises higher levels and, instead, makes programming more complicated. This concept is based on the belief that specifying an operation by selecting it from a list of options, rather than by typing a command or a statement, represents a higher level of abstraction. Most development environments have features based on this delusion. Like the non-procedural languages, creating applications by selecting things from lists is seen as a novel, high-level concept: all we do now, it seems, is tell the computer what we need, and it automatically generates pieces of software for us, even the entire application.

In reality, whether we select options or write statements, our starting elements must be a combination of high-level operations and low-level ones. Thus, even when we communicate with the system by selecting built-in operations, we must create the application's structures – its unique processes, or aspects – and the links between structures. For, if this were not the case, if our contribution were limited to making selections, the only applications we could create would be random and useless combinations of built-in operations.

With a traditional language, we tell the computer what to do by formulating statements, definitions, and expressions. With the new concept, we are shown lists of options, and options within options, and we tell the computer what to

do by selecting entries from these lists. This creates the illusion that we are not programming, that all we must know is how to *select* things. What we must know, though, is the same as before; only the way we apply this knowledge is different.

In principle, one can specify anything by selecting options, but only with trivial requirements is this method more expedient than typing statements. The devices are promoted, however, for *all* applications. Clearly, no one would adopt them if told that they are meant only for novices, or only for solving simple and isolated problems. Thus, the devices must maintain the illusion that, no matter how complex the requirements, all we ever do is make selections; and this is why they end up making programming *more* difficult. But if we believe that one can accomplish more with these devices than without them, we will agree to perform any acts, no matter how illogical, just so that we can use them.

We encounter this delusion, for instance, in the development environments called *visual*, and in those called *by example*. Thus, the concept *query by example* claims to give users the means to perform certain database operations without programming. The concept sounds as if, instead of formulating queries, all we had to do now is show the system some examples of what we need. In reality, since there is no way for a database system to know what we need without being given precise information, we must provide the same specifications as before; and, because we wanted to avoid the traditional method of formulating queries, we end up with a more complicated one.

Thus, to tell the system which records to read, instead of expressing through a traditional language a condition like “products with price less than 100 and quantity in stock greater than 10,” we must perform a number of selections: we select from a list of files “product,” then from a list of fields “price,” then from a list of relational operators “less,” then the value 100, then three more selections for “quantity,” “greater,” and 10, and finally, from a list of logical operators, “and.” Even such trivial acts as the entry of a numeric value like 100 can be reduced to a process of selections: we have all seen systems where a number is displayed for us, and we are expected to increment or decrement it with the mouse until it reaches the desired value. This method takes longer than simply typing the value, but it is an important part in the delusion of high levels: we are now only *selecting* a value, not *specifying* it.

It ought to be obvious that in order to select the right field, operation, or value we must know what these notions are, must appreciate the consequences of each selection and of our particular combination of selections, must understand the significance of operations like “less” or “and” when applied to database fields and records, and so on. Also, the query is meaningless as an isolated function; it is part of an application, so we must be aware at the same

time of the application's other aspects, and of the other uses of those files, records, and fields. In other words, to select the right things we must deal with details and with interacting structures, so the new method does not represent a higher level of abstraction: we must have almost the same programming skills as when specifying those things with statements.

The knowledge that is no longer required – remembering what operations are available, for instance, or the correct format of a statement – is the *easy* part of programming, the *mechanistic* knowledge. These devices impress ignorant practitioners, who lack even this basic knowledge (and are unaware of the required *complex* knowledge lying beyond it), and who, therefore, believe that a substitute for *it* is all they will ever need in order to create applications. Experienced programmers refuse to use these devices, not because they cling to the old methods, as the propaganda tells us, but because they recognize how insignificant their benefits are.

The devices, thus, introduce elaborate procedures as a substitute for the *simple* knowledge involved in programming, but they cannot replace the difficult, *complex* knowledge, which can only develop through personal experience. The immensity of the environment, and the endless novelties that must be assimilated in order to use it, mask the fact that it is still our own skills, not the device, that solve the difficult programming problems. In the end, all the device does is sit between us and our applications, forcing us to express our requirements in more complicated ways than we would through traditional programming. Moreover, despite its low-level features, the device still prevents us from implementing all conceivable alternatives.

4

The most flagrant manifestation of software mechanism is the obsession with ways to avoid programming. Serious programming is indeed a difficult pursuit, but so are other professions. And it is only in programming that the main preoccupation of practitioners has become the *avoidance* of the knowledge, skills, and activities that define their profession. The ignorance pervading the world of programming is so great that the obsession with ways to avoid programming forms its very ideology. The irrationality of this obsession can be observed in this strange phenomenon: as programmers and managers are taught that programming must be avoided at all costs, they end up accepting with enthusiasm any theory or system that claims to eliminate the need for programming, even when this makes application development *more* difficult.

Were they not blinded by their mechanistic delusions, software practitioners would easily recognize that the programming aids are only replacing their

simple, mechanistic activities, and that successful application development entails *non-mechanistic* knowledge. One can attain non-mechanistic knowledge only through personal experience. Thus, as long as they are guided by mechanistic beliefs and seek progress through programming substitutes, the software practitioners deprive themselves of the opportunity to gain this experience. They are trapped, therefore, in a vicious circle: the only knowledge they believe to be required is the mechanistic knowledge they are trying to replace with devices; consequently, they interpret each disappointment, not as evidence of the need for additional, non-mechanistic knowledge, but as a shortcoming of the particular device they are using; so, instead of gaining the additional knowledge through programming, they merely look for another device, and repeat the whole process in a slightly different way.

It is worth repeating these facts, because they are perhaps not as obvious as they appear here. How else can we explain the failure of society to notice the incompetence of our programmers? Endless justifications are being suggested to explain why we must disregard, in the case of programmers, notions that we accept implicitly in any other profession; particularly, the need for personal experience in the tasks defining the profession. For programmers, we have redefined the idea of experience to mean experience in using substitutes for experience.

And so it is how the delusion of software mechanism has given rise to that famous phrase, “without writing a single line of code.” When referring to a programming substitute, this phrase is a promise; namely, that the device will permit us to create applications, or pieces of applications, without any programming. This promise is seen as the most desirable quality of a software device, and software companies will do almost anything in order to realize it – even invent, as we saw previously, devices that make application development more difficult. What matters is only the claim that we no longer have to “write code” (write, that is, statements or instructions).

It is not surprising, of course, to see this phrase employed for devices addressing software *users* – office workers, managers, amateur developers, and the like. Since no device can allow someone without programming knowledge to perform tasks requiring programming, the claim is a fraud. But we can understand the *wish* of naive people to have such a device, and consequently their exploitation by charlatans. What is surprising is to see the same phrase employed for devices addressing *programmers* – those individuals whom one would expect to possess programming expertise (and hence to have no use for such devices), to be proud of their programming capabilities, and even to enjoy programming.

The fact that the software charlatans employ the same means of deception in both cases ought to draw attention to the absurdity of our software culture:

individuals whom we all consider professional programmers have in reality about the same knowledge, ambitions, and expectations as average computer users; like mere users, their chief preoccupation is to improve, not their *programming* skills, but their skills in *avoiding* programming.

And it is not just the software companies that foster these delusions. Researchers in universities participate by inventing mechanistic software theories, the business media by promoting worthless software concepts, corporations by employing programmers who rely on aids and substitutes, governments by permitting the software bureaucracy to exploit society, and in the end, each one of us by accepting this corruption. For, simply by doing nothing, by continuing to worship the software elites and to depend on the software bureaucrats, we are in effect supporting them. The cost of the mechanistic software delusions (probably exceeding one trillion dollars a year globally) is passed in the end to society, to all of us. So, just by doing nothing, we are in effect paying them, each one of us, thousands of dollars every year, and helping them in this way to increase their domination.

5

In chapter 4 we discussed Jonathan Swift's criticism of the mechanistic ideology that was sweeping the scientific world at the end of the seventeenth century; in particular, his attack on the mechanistic language theories (see pp. 315–316). The idea that there is a one-to-one correspondence between language and knowledge, and the idea that languages can be studied, designed, and improved as we do machines, were seen in Swift's time as a foregone conclusion, and were defended by pointing to the successes of mechanism in the natural sciences. Thus, even though the mechanistic theories of language were mere speculations, most scientists were taking them seriously. To ridicule these beliefs, Swift has his hero, Gulliver, describe to us the language machine invented by a professor at the Grand Academy of Lagado.¹

The machine is a mechanical device that contains all the words of the English language, and their inflections. By manipulating a number of cranks, the operator can instruct the machine to generate random combinations of words. And by selecting those combinations that constitute valid phrases and sentences, the professor explains, any person intelligent enough to operate the machine – intelligent enough, that is, to turn the cranks – can produce any text in a particular field of knowledge. Thus, a person with no knowledge of

¹ Jonathan Swift, *Gulliver's Travels and Other Writings* (New York: Bantam Books, 1981), pp. 180–183.

philosophy, or history, or law, or mathematics, can now write entire books on these subjects simply by operating the machine.

The professor emphasizes that his invention is not meant to help a person acquire new knowledge, but on the contrary, to enable “the most ignorant person” to write in any field “without the least assistance from genius or study.”² The machine, thus, will allow an ignorant person to generate any text without having to know anything he does not already know. And this is possible because the person will generate the text (as we say today in programming) “without writing a single line.”

Now, one could certainly build such a machine, even with the mechanical means available in the seventeenth century. Swift is not mocking the technical aspects of the project, but the belief that the difficulty of developing ideas is the mechanical difficulty of combining words. If we hold this belief, we will inevitably conclude that a machine that helps us to manipulate words will permit us to perform the same tasks as individuals who possess knowledge, talent, and experience.

It is obvious that the quality of the discourse generated by a language machine depends entirely on the knowledge of the operator. The machine can indeed produce any text and any ideas, but only by randomly generating all possible combinations of words. So, in the end, it is still the human operator that must decide which combinations constitute intelligent sentences and ideas. Although it appears that the machine is doing all the work and the person is merely operating it, in reality the machine is replacing only the *mechanical* aspects of language and creativity.

Thus, a person using the machine will not accomplish anything that he could not accomplish on his own, simply by writing. Now, however, since he is only *selecting* things, it can be said that he is generating ideas “without writing a single line.” Whatever the level of intelligence of a person, it is in fact more difficult to generate a piece of text by operating this machine than by directly writing the text. But if we believe that it is the *mechanical* acts involved in writing that make writing difficult, or if we have to employ as writers individuals known to be incapable of writing, we might just decide that language machines make sense.

Returning to our software delusions, we indeed believe that the difficulty of programming lies in its *mechanical* aspects, in combining pieces of software; and, what is worse, we indeed have to employ as programmers individuals known to be incapable of programming. So we have decided that *programming* machines make sense.

The similarity between Swift’s hypothetical language aid and our real

² Ibid., p. 181.

programming aids is striking. We note, in both cases, devices that address ignorant people; assure them that they don't need to know anything they don't already know; promise them the power to perform tasks that require, in fact, much knowledge; and reduce their involvement to a series of selections.

The similarity is not accidental, of course. We already know that our software delusions and our language delusions stem from the same belief; namely, the belief that the elements of software structures and language structures correspond on a one-to-one basis to the elements that make up reality. So we must not be surprised that devices based on *software* delusions end up just like the device invented by a satirist to mock the *language* delusions.

Swift was trying to demonstrate the absurdity of the mechanistic language theories by exposing their connection to the belief that a mechanical device can replace human knowledge. But today, through the mechanistic *software* theories, we are actually attempting to realize this fantasy: we are building *software* devices to replace human knowledge – programming knowledge, in particular. Concepts that were only academic speculations in Swift's time, easily ridiculed, have become a reality in our time in the world of software. The kind of device that three centuries ago was only a fantasy – a satirical exaggeration of a delusion – is actually being built today by software companies, and is being used by millions of people.



There is no better way to illustrate the essence of software charlatanism than by imagining how the professor from Lagado would design his language machine today. He would make it a software device, of course, rather than a mechanical one. And, as a matter of fact, it is quite easy to design a software system that allows anyone – including persons who are normally unable to express themselves – to produce books in any domain “without writing a single line.” To imagine this device, all we have to do is combine the concepts implemented in Swift's language machine with those implemented in our software systems.

The promise, thus, would be the familiar claim that the only thing we need to know is how to operate the device, and that this knowledge can be acquired in a short time. To operate the mechanical language machine, all they did was turn cranks; to operate a modern language machine, all we would do is “point and click” with a mouse. We are assured, in both cases, that the power of the device is ours to enjoy “at a reasonable charge, and with a little bodily labour,”³ and only by making selections: we would never have to write a single sentence. The phrase we would use today is “powerful yet easy to use.”

³ Ibid.

Let us examine some of the possibilities. Instead of typing words, we can have the system display them for us in the form of selections within selections. If, for example, we need the sentence “the dog runs,” we first select the grammatical function by clicking on *noun*; this displays a list of noun categories, and we select *animal*; within this category we select *domestic*, and finally *dog*; what is left is to click on *singular* and *definite article*. Then, for “runs” we select the grammatical function *verb*, which displays a list of verb categories; we select *action*, within this category we select *motion*, and finally *run*; we then click on *present tense*, *third person*, and *singular*, and the complete sentence is displayed for us.

The popular expedient of *icons* could be profitably employed to help even illiterate persons to use this system: if tiny pictures were used to depict words, categories, and grammatical functions (a picture of a dog for “dog,” an animal together with a man for “domestic,” a running figure for “run,” one and two objects for “singular” and “plural,” etc.), even those of us who never learned to read and write could benefit from the power of language machines.

It is obvious that, with such a system, anyone could generate text on any subject without writing a single line. And future versions could introduce even more powerful features – built-in sentences, for instance. Instead of words, we would be able to select entire sentences, and even entire paragraphs, from lists of alternatives representing classes and categories of topics. With only a little practice, anyone would then be able to generate page after page of exquisite text just by pointing and clicking.

The more elaborate this imaginary language system becomes, the easier it is to recognize its similarity to our software systems – our programming aids, in particular. But, while few people would be deceived by a language machine, the whole world is being deceived by the software charlatans and their application development machines. Not even illiterates could be persuaded to try a device that promises to replace writing skills. But the most important individuals in society – decision makers working in universities, corporations, and governments – keep trying one software theory after another, and one programming substitute after another, convinced that a device can replace *programming* skills.

6

Whether addressing programmers or software users, an *honest* development system simply provides low-level elements and the means to combine them so as to create the higher levels. The low levels come (for programmers, at least) in the form of general-purpose programming languages; and, when practical,

higher levels are available through existing subroutines. Systems that provide *only* high levels, and claim that it is possible to create any application in this manner, are dishonest: they invariably end up reinstating the low levels in a different, and more complicated, form. These systems are for programming what language machines are for writing: not useful tools, but means of deception and exploitation. Their purpose, we saw, is to induce ignorance and dependence, by consuming our time and preventing us from improving our skills.

Honest systems allow us to create the higher levels on our own, and to select any subroutines we like. With honest systems, therefore, we can choose any combination of low-level elements and built-in operations. Dishonest systems provide an environment with high starting levels, and add the low levels as a special feature. The software charlatans have reversed, in effect, the principles of programming: instead of a simple system based on low levels, where we can create the high levels independently, they give us a complicated environment based on high levels, where the low levels are provided as “enhancements.” What we had all along in any programming language – the low levels – is presented now as a new and powerful feature of their high-level environment. Instead of programming being the standard development method, and the high levels a natural outcome, they make the high levels the standard method, and turn programming into a complicated extension.

Clearly, if we use a general-purpose development system, if we want to create original applications, and if these applications require a particular level of detail and functionality, our lowest-level elements must be the same mixture of variables, conditions, loops, and statements no matter what development method we use.

The software charlatans prefer environments based on high levels because this is how they can induce dependence. A system based on low levels and subroutines leaves us free to develop and maintain our applications in any way we like. The dishonest systems lure us with the promise of high starting levels, but must include the low levels anyway. They lose, therefore, the only benefit they could offer us. But, because we trusted them and based our applications on their high levels, we will now depend forever on them and on the software companies behind them. While no dependence is possible when using traditional development methods, it is quite easy to control our work, our time, our knowledge, and our expectations through systems based on high levels. For, instead of simply developing applications and expanding our programming skills, we are now forced to spend most of our time with the problems generated by the systems themselves, with complicated concepts, with special languages, and with their endless changes.

The only time a high-level system is justified is when its functions cannot

be effectively implemented as subroutines. This is the case, typically, in systems meant for highly specialized applications. Thus, operations involving indexed data files can be added as subroutines to any language. They are more convenient when implemented in the form of statements (as in COBOL), but it would be silly to adopt a new language, or a whole development environment, just for this reason. On the other hand, the features found in an advanced file editing system cannot be simply added to a language as subroutines, because, by its very nature, the editing system must have its own environment (windows, commands, special use of the keyboard, etc.). And what an honest system does, in this case, is make it as easy as possible to transfer the files to and from other systems.



It is worth repeating here that “subroutine” refers to a broad range of high-level software elements, including functions, procedures, subprograms, and the like, which may be explicit or implicit. This term refers, thus, to any elements that can be implemented as a *natural extension* of a general-purpose programming language. The subroutines that perform file operations, for example, may be implemented by way of functions (in a language like C) or as ordinary statements (in a language like COBOL). The important point is that the foundation of the application be a general-purpose language, not the high-level entities of a development environment.

And I refer to individual statements, conditions, iterations, etc., as “low-level” software elements only because they are lower than subroutines, or built-in operations, or the high-level functions provided by development environments. But these “low-level” elements are what we find, in fact, in general-purpose languages (like COBOL and C) called “high-level” (to distinguish them from assembly languages, which use true low-level elements).

This confusion in terminology is due to the software mechanists, who have distorted the meaning of low and high levels by claiming that it is possible to raise forever the level of the starting elements. Thus, the term “fourth generation” (4GL) was coined for the languages provided by development environments, and “third generation” for the traditional high-level languages, in order to make environments look like an inevitable evolution. Assembly languages were declared at the same time to be “second generation,” and machine languages, which use even lower-level elements, “first generation.”

The level of these languages, however, has little to do with an advance in programming concepts. Thus, the first three generations are still in use today, and will continue to be, because their respective levels are the only way to implement certain types of operations. It is true that, historically, we started

with first generation and only later invented the second and third; but this doesn't prove that there can exist general-purpose languages of even higher levels. And it is true that, in most programming tasks, we were able to replace the first two generations with the third without reducing the functionality of the resulting applications; but it doesn't follow that we can repeat this success, that we can develop the same applications starting from even higher levels.

Everyone agrees that it is more efficient to start from higher levels, and that we should use the highest-level entities that are practical in a given situation. But, as we saw earlier, for typical business applications this level cannot be higher than what has been called third generation. While we may agree that the first three generations represent a certain progression in programming concepts, the fourth one is a fraud. Not coincidentally, it was only when the fourth one was introduced that the term "generation" was coined; formerly we simply had "low-level" and "high-level" languages.

The fourth generation is a fraud because it is not, relative to the third, what the third is relative to the second. In the transition from second to third generation, practically all language features became simpler: the use of memory variables, conditions, iterations, subroutine calls, etc., are indeed at a higher level than in assembly languages. In the transition from third to fourth generation, on the other hand, practically all features remained unchanged. What the software charlatans did was merely add a few high-level features (built-in operations for simple reports, user interface, etc.) and bundle the resulting languages with proprietary development environments. A programmer can enjoy the same blend of low and high levels by starting with traditional languages (COBOL, C, etc.) and adding subroutines and similar features, created by himself or by others.

To function as a higher level, a language must provide more than just a few higher-level functions; it must provide a higher level for variables and arrays, for flow-control constructs, for subroutines and their parameters, and so on. It is precisely because no such features can exist (in the case of general-purpose languages and general business applications), that the third-generation features had to be retained. The idea of a fourth generation was merely a way to get ignorant practitioners to abandon the straightforward programming languages, and to depend instead on complicated and expensive development environments.⁴

⁴ In more than forty years of programming – from simple utilities and applications to large data management systems and business systems – I have never encountered a situation where I could benefit from a commercial development environment. Even when the project calls for a higher-level programming method, I find it more expedient to implement my own, simple, customized environment (using "third-generation" and "second-generation" languages) than to depend on those monstrous systems sold by software companies.

7

As an example of development environments, let us examine the communications systems. If what we need in our business applications is high-level operations in the domain of communications (say, transferring data under various protocols between computers, or converting files from one format to another), nothing could be simpler than providing these operations in the form of subroutines. We could then develop the applications in any programming language we like, and invoke these operations simply by specifying a number of parameters.

Needless to say, this is *not* how the popular communications systems make their operations available. What programmers are offered is a whole environment, where the operations are invoked interactively. Then, because the interactive method is impractical when the operations must be part of an application, these systems also provide a “powerful feature”: a programming language. (To further distract us, euphemisms like “scripts,” “macros,” or “command files” are employed to describe the resulting programs.) In short, we are taken back to the lower levels of traditional programming. But we already had programming languages; all we wanted was a few high-level communications operations. Instead, we must get involved with, assimilate, and then become dependent on, yet another system, another language, another software company, and the related documentation, newsletters, seminars, websites, version changes, bug reports, and so on.

Most of these activities are spurious, in that they are caused, not by the communications operations we needed, but by the environment we were forced to adopt in order to have these operations. What is worse, the languages that come with these environments are more primitive and less efficient than the general-purpose languages we already had. Only ignorant programmers can be deceived by this fraud, of course; true professionals recognize that these systems are unnecessary, that their sole purpose is to prevent programming freedom. The popularity of development environments, and the ease with which practitioners can be persuaded to depend on them, demonstrates therefore the incompetence that pervades the world of programming. It is in the interest of the software companies to maintain this incompetence. Thus, by providing environments instead of honest development systems, they ensure that programmers waste their time with spurious activities. For, with proprietary environments we only acquire some narrow skills; with general-purpose systems and languages, on the other hand, we gain true expertise – knowledge and experience that will benefit us in all future programming work.

Communications systems are only one kind of environment, of course. If we are to depend on development environments for our high-level operations, we will also need systems for display, for user interface, for database operations, for graphics, for reporting, for system management, etc. – each one with its own language, documentation, newsletters, seminars, changes, bugs, and so on.⁵

Development environments must include programming languages because their high-level operations, no matter how impressive they may be on their own, are only useful when combined with *other* operations. An application is not simply a series of high-level operations. The operations provided by one system are related to those provided by another, and also to the operations developed specifically for that application. The relations between these operations occur mainly at low levels, so they must be implemented through conditions, loops, statements, and variables; in other words, through the same low-level elements as those found in the traditional programming languages. Like the language machine we examined previously, the environments promise us high levels, but provide in reality the same mixture of levels we had all along. To develop a given application we need the same knowledge as before, but applying that knowledge is now much more difficult.

The complications created by this charlatanism are so great that a new kind of system had to be invented, whose only purpose is to help programmers and users connect the operations of the other systems, or transfer data from one system to another; its only purpose, thus, is to solve the problems created by the idea of software environments. These new systems come, of course, with their own environments, languages, procedures, documentation, newsletters, seminars, bugs, changes, and so on. Another kind of system engendered by software charlatanism is the one meant to standardize the operations provided by other systems – to sit above them, as it were, and make their diverse operations available in a common format. Every software company tries to establish *its* system as the standard one, but this struggle merely results in even more facts, languages, procedures, documentation, reviews, etc., that programmers must assimilate.

These complications, to repeat, are a result of the reversal of programming principles: instead of starting with low-level elements and creating the higher levels freely, programmers are forced to develop applications starting with

⁵ Thus, software reseller Programmer's Paradise boasts on its catalogue cover, "20,000+ software development tools" (for example, the issue Nov–Dec 2008). Perhaps 1 percent of them are genuine programming tools. The rest are environments and the endless aids needed to deal with the problems created by environments. Individuals who need such tools are not true programmers, but a kind of users: just as there are users of accounting systems and inventory management systems, *they* are users of development systems.

high-level elements. The low levels are then provided only *through* the development environments, and *through* the high levels, thus establishing the dependence.

A system based on low levels and subroutines *also* offers the benefits of high-level elements, and without inducing any dependence. After all, we already have many programming languages; and through these languages, we can create software levels that are as low or as high as we want. Software companies do not promote environments because our general-purpose languages are inadequate, but because traditional concepts like subroutines do not allow them to control our work and our applications as do these environments.

The Delusion of Methodologies

1

So far we have discussed the development of applications mainly from the perspective of programmers. Let us see now how the mechanistic delusions affect the expectations of the *users* of applications – those individuals whose needs are to be embodied in the new software.

When developing a new application, managers familiar with the relevant business practices cooperate with analysts and programmers. The resulting software, thus, will reflect not only programming skills, but also the knowledge and experience of users. And the mechanistic theories and methodologies expect these individuals to express their knowledge and their requirements precisely, completely, and unambiguously; that is, to reduce knowledge and requirements to a form that can be used by analysts and programmers to develop the application. We will examine this absurdity in a moment, but first let us briefly discuss the alternative.

Instead of developing custom software, users can procure ready-made (or what is known as packaged, or canned) applications. With this alternative, the application is available immediately, thus bypassing the lengthy and difficult stages of design and programming. From what we have already discussed, though, it should be obvious that packaged applications are part of the same delusion as all ready-made, or built-in, pieces of software: the delusion of high levels, the belief that one can accomplish the same tasks by starting from high-level software elements as when starting from low-level ones. This delusion finds its ultimate expression in the idea of ready-made applications: the starting level is then the top element itself, the complete application, and the impoverishment is total. From the infinity of alternatives possible for the top

element when *programming* the application, we are now left with only one alternative: the particular combination of operations built into the package by its designers. Most packages include options, of course, for some of their built-in operations. But combinations of options still provide only a fraction of the combinations of operations that may be required by an organization. So, in the end, packages remain a poor substitute for custom applications.

Organizations are tempted by the promise of packaged applications because they underestimate the limitations they will face later, when they get to depend on this kind of software. And even when the users realize that the package permits them to implement only *some* operations, and addresses only *some* of their needs, they still fail to appreciate the real consequences of inflexible software. What they tend to forget is that their needs and practices evolve continually, so their software applications must evolve too. It is difficult enough to judge whether a certain application can answer our *current* needs (the only way to be absolutely sure is by running it live, by *depending* on it); but it is impossible to assess its usefulness for the next ten years, simply because we cannot know *what* our needs will be. No one can predict the changes that an organization will face in the future. How, then, can anyone expect a piece of software that is based on a particular combination of built-in processes and operations to cope with such changes?

Note that it is not the *quality* of the application that is at issue here: no matter how good and useful it is today, and even if the company supporting it will bring it up to date regularly in the future, it will always be a *generic* piece of software, designed to answer only that subset of needs common to many organizations; it cannot possibly adapt to the specific needs of every one of them.

It is quite incredible, thus, to see so many organizations depend on packaged software; and they do, not just for minor applications, but also for their important business needs. Most packages fail, of course, so we must not be surprised at the frequency with which these organizations try new ones. The failure of a package rarely manifests itself as major deficiencies, or as software defects. What we see typically is a failure to answer the needs of its users, something that may only become evident months or years after its adoption. Since this type of failure is so common, the reason why organizations continue to depend on packages is, clearly, not their usefulness, but the incompetence of the software practitioners: if programmers lack the skills necessary to create and maintain applications, ready-made software, however unsatisfactory, becomes an acceptable expedient.

More subtle and more harmful than the inadequacy of an application is the tendency of users to lower their expectations in order to match its limitations. In other words, instead of rejecting an inadequate application, they modify the

way they conduct their affairs so as to be able to use it. To help them rationalize this decision, the software elites stress the importance of adopting the latest “technologies” – relational databases, object-oriented environments, graphic user interface, client-server systems, and so forth. Ignorant users are impressed and intimidated by these concepts, so they end up interpreting the application’s shortcomings as modern and sophisticated features which they don’t yet appreciate. Thus, instead of objectively assessing the application’s usefulness, they merely judge it by how closely it adheres to the software innovations promoted by the elites, even if these innovations are worthless. So, in the end, the application appears indeed to satisfy their requirements; but this is because they agreed to replace their true requirements with spurious ones.

2

Having established that packages are rarely a practical alternative for serious applications, let us return to the subject of software *development*. Developing their own applications is what many organizations must do, even if lacking the necessary skills, because this is the only way to have adequate software.

An application, we recall, consists of many structures, all sharing the same software entities (see “Software Structures” in chapter 4). These structures are the various *aspects* of the application – the processes implemented in it. Each structure, thus, is one way of viewing the application; and it is this system of interacting structures that constitutes the actual application. Although in our imagination we can treat each aspect as a separate structure, the only way to create the application is by dealing with several structures at the same time. This is true because most entities in a piece of software – most statements and modules – are affected by several aspects of the application, not just one. When writing a statement, for example, it is seldom sufficient to think of only *one* logical structure; we may well perceive a particular structure as the most important, but the same statement is usually an element in other structures too. It is this capability of software entities to be part of several structures simultaneously, and hence link them, that allows software applications to mirror our affairs. This capability is important because our affairs consist of processes and events that already form interacting structures.

If this is what software applications actually are, let us review what the software theories assume them to be. Applications, the theories tell us, must be developed following a *methodology*. Although many methodologies have been proposed, all are ultimately founded on the same fallacy; namely, the belief that it is possible to reduce a software application to a *definition*. The definition of an application is a set of specifications (formal descriptions, flowcharts, block

diagrams, and the like) believed to represent, *precisely and completely*, the actual software. Methodologies, thus, are a manifestation of the mechanistic belief – the belief that a complex structure (the software application, in this case) can be reduced to simple ones.

To define an application, users and analysts spend many hours discussing the requirements – the business practices that are to be embodied in the application. This activity is known as analysis and design, and the methodologies prescribe various steps, which, if rigorously followed, are said to result in a complete definition; namely, a definition that represents the application as precisely as drawings and specifications represent a house or a car. It is believed, thus, that a set of mechanistic expedients can capture all the knowledge inhering in a complex phenomenon: the structures that make up the application, their interactions, and their effects when the application is running.

The reason we start with a definition, of course, is that we prefer to work with specifications rather than the statements of a programming language. Deficiencies, for example, are easier to correct by modifying the definition than by rewriting software. Thus, we are told, if we follow the methodology, we should be able to create the entire application in the form of a definition, and then simply translate the definition into a programming language. To put this differently, the methodologies claim that it is possible to represent an application with expedients other than the software itself – expedients that are simpler than software, and accessible to users and programmers alike. Although simpler than the actual application, these expedients represent it precisely and completely. The definition *is*, in effect, the application.

The fallacy of this claim ought to be obvious: if it were possible to express by means of diagrams, flowcharts, etc., all the details of the application, we wouldn't need programming languages. For, a compiler could then translate the definition itself into the machine language, and we wouldn't need to write the programs. In reality, definitions are simpler than programs precisely because they do *not* include all the details that the programs ultimately will.

So definitions are *inaccurate* representations of the application. They are useful only because people can interpret them, because people can add some of their own knowledge when converting them into software. One reason why definitions are simpler than programs, thus, is that they need not be perfect. An error in the program can render the application useless, but in the definition it is harmless, and may even go unnoticed. The impossibility of translating automatically definitions into software proves that definitions are incomplete, faulty, and ambiguous, and require human minds to interpret and correct them.

But an even more important reason why definitions are simpler than

programs is that they represent *separately* the software structures that make up the application. The difficulty in programming, we saw, is dealing with several structures simultaneously. Our programming languages permit us to create software entities that can be shared by diverse structures, and this is why it is possible to develop useful applications. In a definition, on the other hand, we usually specify each structure separately: the business practices, the database relations and operations, the display and report layouts – we strive to represent each one of these processes clearly, so we separate them. Even if we *wanted* to relate them in the definition it would be difficult, because the diagrams, flowcharts, and descriptions we use in definitions are not as versatile as programming languages. Definitions are simpler than programs, thus, because most specifications do not share their elements, as software structures do. What this means is that a definition *cannot* represent the application precisely and completely. So the methodologies are wrong when claiming that definitions are important.

The fallacy of definitions is easy to understand if we recall the concept of simple and complex structures. A definition is, in effect, the reification of a complex structure (the application) into its constituent simple structures. It is, thus, an attempt to reduce a complex phenomenon to a mechanistic representation. This can be done, as we know, only when the separated structures can usefully approximate the actual phenomenon. In the case of software phenomena, this can be done for trivial requirements. For typical business applications, however, mechanistic approximations are rarely accurate enough to be useful. In the end, we like software definitions for the same reason we like all other mechanistic concepts: because of their promise to reduce complex problems to simple ones. Definitions are indeed simpler than the applications they represent, but they are simpler because they are only approximations.

Thus, since applications cannot be represented accurately by any means other than the programs themselves, the conclusion must be that definitions are generally irrelevant to application development. They may have their uses, but their importance is overrated. No definition can be complete and accurate, and an application created strictly from a definition is useless. Application development cannot be reduced to a formal activity, as the software theorists say. Since no one can specify or even envisage all the details, and since most details will change anyway (both before and after the application is completed), it is futile to seek a perfect set of specifications. Some brief and informal discussions with the users are all that an experienced programmer needs in order to develop and maintain an application.



The failure of the mechanistic concepts in the early days were so blatant that the software gurus had to modify their methodologies again and again. The invention of new methodologies, thus, became a regular spectacle in the world of programming, and there were eventually nearly as many methodologies as there were gurus. (Most methodologies are known by the name of their creators, a practice borrowed apparently from the world of fashion design.)

Some methodologies tried to eliminate the rigidity of the traditional development phases, and introduced notions like prototyping and stepwise refinements; others attempted to modify the traditional roles played by users, analysts, and programmers. But, in the end, no matter how different they may appear to the casual observer, all methodologies are alike. And they are alike because they all suffer from the same fallacy: the belief that indeterministic phenomena – the applications, and their development and use – can be treated as mechanistic processes. The idea of methodologies, thus, is just another manifestation of the belief that programming expertise can be replaced with some easy skills – the skills needed to follow rules and methods.

The similarity between the various methodologies is betrayed by the trivial innovations their creators introduce in an effort to differentiate themselves. For example, they use pretentious terms to describe what are in fact ordinary features, in order to make these features look like major advances. But most ludicrous is their preoccupation with the graphic symbols employed in diagrams, as if the depiction of processes, operations, and conditions with one symbol rather than another could significantly alter the outcome of a development project. For example, the traditional rectangular boxes are replaced with ovals, or with a shape resembling a cloud, or a bubble, or one known as a *bubtangle* (a rectangle with rounded corners). And we must remember that these idiocies are discussed with great seriousness in books and periodicals, and are taught in expensive courses attended by managers and analysts from the world's most prestigious corporations.

Programming methodologies, thus, are like the development environments we discussed previously: they provide elaborate systems to replace the *easy* aspects of programming, those parts demanding mechanistic knowledge; but they cannot replace what are the most important and the most difficult aspects, those parts demanding complex knowledge. Since the same knowledge is required of people to create a serious application whether or not they use a methodology, the methodologies, like the development environments, are in the end a fraud. They are another form of software exploitation, another way for the software elites to prevent expertise and to induce dependence on systems and devices which they control.

When a methodology appears successful, its contribution was in fact insignificant. For, why should some techniques that work for one organization

fail to work for others? It is the people, obviously, that made the difference. When people have the necessary knowledge, they will develop applications with or without a methodology; and when they lack this knowledge, no methodology can help them. Development environments, we saw, promise programmers and users higher levels of abstraction, and then trick them into working at low levels, as before. Similarly, methodologies promise them simpler, high-level concepts, and then demand the same skills as before. In both cases, this charlatanism complicates the development process, so inexperienced practitioners are even less likely to succeed. Besides, they waste their time now assimilating worthless concepts, instead of using it to improve their skills by creating and maintaining applications.

3

The delusion of methodologies and definitions is reflected in the distorted attitude that everyone has toward the subject of *maintenance*. Software maintenance is the ongoing programming work needed to keep an application up to date. And all studies agree that, for most business applications, this work over the years exceeds by far the work that went into the initial development. We should expect the theorists, therefore, to propose more solutions to the problems arising in maintenance than to those arising during development. What we find, though, is the exact opposite: all theories and methodologies deal with the creation of new applications, and barely mention the subject of maintenance. Moreover, we find the same distorted attitude among corporate managers: maintenance is treated as incidental work, is avoided whenever possible, and is relegated to the least experienced programmers.

In reality, the obsession with new applications is a reaction to the problem of programming incompetence: because programmers cannot keep the existing applications up to date, new ones must be developed. But without proper maintenance the new ones quickly fall behind, so the users find themselves in the same situation as before. At any given time, then, companies are either installing new applications, or struggling with the current ones and looking forward to replacing them. The software elites encourage this attitude, of course, as it enhances their domination. They present the latest fads – fourth-generation or object-oriented systems, CASE tools or relational databases, graphic interface or distributed computing – as revolutionary advances, and as solutions to the current problems. Their applications are inadequate, the companies are told, because based on old-fashioned software concepts. They must replace them with new ones, based on these advances.

So the preoccupation with new applications helps everyone to rationalize

the failure of maintenance. It takes great skills to modify a live application quickly and reliably. In contrast, creating a new application from a definition, as the methodologies recommend, is relatively easy. It is easy because the neat definition is only a simplified version of the actual application. As we saw, definitions can only *approximate* the true, complex needs. But the belief that the next application will be perfect inspires everyone with confidence. So a new development project, using the latest development fads, always looks like a wise decision.

To put this differently, practitioners prefer a new application to maintenance because new projects make self-deception possible. A methodology permits them to create, instead of the *required* application, an imaginary, simpler one: the application matching a neat definition and their limited skills. And when *that* application proves to be inadequate, the practitioners still do not suspect their practices. They blame the changing requirements, or the imperfection of the original specifications. They refuse to see these facts as a reality they must cope with, as the very essence of business software. So, instead of accepting the facts, they continue to claim that their practices are sound, and that precise definitions are possible. In other words, if reality does not match the mechanistic software principles, something is wrong with reality.

In new development projects, then, self-deception helps practitioners to deny their failures and to cling to the easy, mechanistic concepts. And they dislike maintenance because, in this type of work, self-deception cannot help them. Each maintenance project is relatively small and well-defined, so it is harder to replace it with an imaginary, simpler one. Ultimately, in maintenance work it is harder to find excuses for failures.



We note a marked discrepancy between the *perception* and the *reality* of applications. On the one hand, everyone strives to create a perfect application – by following a strict methodology, and by using the latest development systems. It is far more expensive to modify the software itself later, we are told, so we must eliminate the imperfections in the design stage. This is why definitions are important. On the other hand, all studies show that less than 5 percent of new applications are adequate. The others must be modified if they are to be used at all, and many are so different from the actual requirements that they must be abandoned. Moreover, even those that are adequate must immediately start a process of ongoing modifications, simply because business requirements change constantly.

Thus, whether it is the original differences (due largely to the fact that no definition can reflect the actual requirements) or the future ones (due to the

normal, unpredictable changes in requirements), it is obvious that modifying business applications is an essential programming activity. Yet, for more than forty years, all theories and methodologies have been attempting to create “perfect” applications; that is, applications matching some fixed specifications, and requiring as few changes as possible. In reality, all software changes are alike – whether due to faulty specifications, or varying user preferences, or the need for additional features, or the adoption of new business rules, or some external factors. So, if we must be able to deal with endless changes in any case, the idea of a perfect application is meaningless, and there is no point in trying to design one initially.

It is wrong, in fact, even to think of maintenance as modifying the application. The role of business software is to satisfy, at any given time, the current needs. An application, therefore, must be seen as that particular software system which accomplishes this. Business needs change constantly, so the application must change too. Thus, rather than first developing an application and then maintaining it, it is better to think of this work as a continuous, never-ending development.



We find further evidence of the distorted attitude toward maintenance in the notion of application *life cycles*. All experts agree that applications cannot last more than a few years. So, even while encouraging us to create a new one, they warn us to prepare for its demise. Borrowed from biology, the idea of life cycles holds that software resembles live things, so the existence of an application can be divided into stages: birth (definition of requirements), growth (development and testing), maturity (normal operation), and death (obsolescence). Each application represents a cycle, and is followed by another one, and then another one, forever.

But this is an absurd idea, contrived specifically in order to justify the need for new applications. Software, by its very nature, is modifiable. In principle, then, an application never needs to be replaced; it only needs to be kept up to date. Everyone acknowledges the need for changes, and acknowledges also the inability of programmers to implement them. So the idea of life cycles was introduced as a compromise: every few years, a new application is created in order to implement *together* all the changes that should have been implemented *one at a time* in the past. The theorists and the practitioners can now defend the lack of proper maintenance, and hence the need for a new application, by invoking the idea of software life cycles. This logic, however, is circular; for, the idea of life cycles was itself an invention, a response to the incompetence that prevents proper, *ongoing* maintenance.

Instead of trying to eradicate the incompetence, everyone looks for ways to rationalize it.

Business software can fulfil its promise only if it is as changeable as the business issues themselves: inflexible business software can be as bad as inflexible business practices. Thus, replacing the whole application from time to time is a poor substitute for the ability to satisfy new needs as soon as they arise. So the ultimate price we pay for distorting the subject of maintenance is having to depend on perpetually inadequate applications. This is true because, even though an inadequate application is eventually replaced, it reaches that condition gradually, one unsatisfied requirement at a time. This means that it was *always* inadequate, even in its period of normal use. The difference between that period and the time when it is actually replaced is only in the *degree* of inadequacy; namely, how far it is from the users' actual needs, how many unsatisfied requirements have accumulated to date.¹

4

The delusion of methodologies and definitions is also demonstrated by the failure of CASE (Computer-Aided Software Engineering, see pp. 521–522). The elimination of programming from the process of application development was seen by most theorists as the undisputed next step in development tools, as the ultimate benefit of software engineering. Ambitious CASE systems were promoted for a number of years with the claim that managers and analysts could now create directly, without programming, applications of any complexity – simply by manipulating block diagrams, flowcharts, and the like, on a computer display. The system would guide them in creating the definition, and would then translate the definition automatically into software.

The belief that an application can be generated automatically is a logical consequence of the belief that a definition can represent all the knowledge embodied in an application. (Could definitions do that, automatic programming would indeed be possible.) The CASE fantasy, thus, was born from the

¹ A properly maintained application never needs to be replaced, because it always has what the users need. The longest I maintained one of my applications is thirty-one years (until the manufacturing company using it ceased production). This was a complex, integrated business system, which combined all the computing needs of that company. At any given time there was a list of requirements, some of them urgent; but I always implemented them, so no one ever saw the need for new applications. The system kept growing, and was eventually a hundred times larger than it had been in the first years, due to countless new functions; but no one perceived these developments as new applications. Most work, though, was in modifying existing parts (replacing or adding features and details).

concepts of methodologies and definitions that we have just discussed – concepts which *continue* to dominate the programming theories, despite the failure of CASE. No one seems to realize that, if CASE evolved from these concepts, its failure proves the fallaciousness of these concepts too. Let us analyze this connection.

Even when following a methodology, people do more than implement rules and standards. The software created by programmers contains more than what the analysts specified in their definition, and the definition created by analysts contains more than what the users specified in their requirements. Each individual involved in the development of the application has the opportunity to add some personal knowledge to the project, but this is largely an unconscious act. Simply to *understand* a set of requirements or specifications, the person must *interpret* them; that is, he must combine the knowledge found in the document with some previous knowledge, present in his own mind. For, if this were not the case, if the only thing that analysts and programmers did were follow rules and methods, then a person who knows nothing about software or about a particular company, but who can follow rules and methods, could also develop applications.

The knowledge missing from the formal requirements and specifications, and hence contributed by individuals, varies from general facts on computers and software to details specific to their organization, and from common business practices to the knowledge shared by people living in a particular society. It is precisely because most people already possess this kind of knowledge that we take it for granted and do not include it in instructions and documents. Recall also that the most important part contributed by human minds constitutes *non-mechanistic* knowledge: not isolated knowledge structures, but the complex structure that is their totality. The capacity for non-mechanistic knowledge must be provided by human minds because it cannot exist in simple structures like instructions or diagrams.

Thus, all the people involved in the development of an application may be convinced that they are following the rules prescribed by the methodology, while depending on personal knowledge and experience to fill in the missing pieces, or to resolve the ambiguities and inconsistencies found in specifications. If the application is successful, they will praise the methodology, convinced that it was the principles of software engineering that led to their success. Most likely, they will not realize that it was in fact their own minds that provided the most important part (the non-mechanistic knowledge), and that the principles, theories, and methods addressed only the simple part (the mechanistic aspects of the project).

Clearly, if the methodology provides only mechanistic principles while our activities are mostly non-mechanistic, the only way to use a methodology is by

taking its practical parts and ignoring or overriding the rest. People may be convinced that they are *following* the methodology, when they are using it *selectively*. So it is not too much to say that, to develop an application successfully, people must work *against* the methodology: if they rigorously followed the mechanistic principles, they would never complete the application. Thus, when a software project is successful, this is not *due* to the methodology but *despite* it.

And it is during programming that people make the greatest contribution. For it is in programming, more than in any other activity, that people have the need and the opportunity to override the rules imposed by a mechanistic methodology. So it is the programmers – more than the managers with their specifications, or the analysts with their definitions – that must use the non-mechanistic capabilities of their minds. We can perhaps delude ourselves in the early stages of development that specifications and definitions represent the application completely and precisely. But if we want to have a useful application, we must permit human minds to deal at some point with the missing pieces, with the ambiguities, and with the inconsistencies. It is during programming, therefore – when the application is created and tested, when it must mirror reality if it is to be used at all – that the delusions of formal methodologies and precise definitions, of neat diagrams and flowcharts, must come to an end.



It should be obvious, then, why CASE failed. The CASE systems were based on methodologies: they literally incorporated some of the popular methodologies, thus allowing managers and analysts who wished to follow a particular methodology to do so through a software system rather than on their own. The system could now *force* people to follow the methodology, eliminating the temptation to omit or modify some of the steps – what was believed to be the chief cause of development failures. Since the methodology was now part of the development environment, the experts claimed, anyone could enjoy its benefits; and since the resulting specifications and definitions were stored in the computer, the system could use them to generate the application automatically, eliminating the programming phase altogether.

CASE failed because it eliminated the opportunities that people had to *override* the methodologies and definitions. By automating the development process, CASE made it impossible for people to contribute any knowledge that conflicted with the mechanistic software theories. They could only use now trivial, mechanistic knowledge, which is insufficient for developing serious applications. What CASE eliminated – what the software mechanists thought

was the cause of development failures – was in fact the very reason why methodologies and definitions appeared occasionally to work: the contribution made by people when, out of frustration, and perhaps unconsciously, they were using their non-mechanistic capabilities to override the methods, rules, and specifications. Thus, the failure of CASE proves that people normally contribute to the development process a kind of knowledge – non-mechanistic knowledge – that cannot be replaced with formal methodologies and theories.

There is another way to look at this. A CASE environment is logically equivalent to a traditional development environment where the users, the analysts, and the programmers follow a methodology *rigorously*; where analysis and design, specifications and definitions, theories of programming and testing, are all implemented exactly as dictated by the principles of software engineering; where everyone refrains from *interpreting* the specifications or the definitions; where no one uses personal knowledge to add details to the formal documents, or to resolve ambiguities and inconsistencies. A CASE environment is equivalent to all this because, when the methodologies and programming theories are part of the development system, people are *forced* to follow them rigorously.

Logically, then, the only difference between a CASE environment and a traditional environment is the non-mechanistic knowledge contributed by people – the knowledge that *cannot* be incorporated in a CASE system. So, if CASE failed, we must conclude that this knowledge plays a critical part in a development project. With traditional development methods, when people possess this knowledge the project is successful, and when they do not the project fails. In a CASE environment, people had no opportunity to use this knowledge, whether they possessed it or not; so the result was the same as when people used traditional development methods *and* lacked this knowledge. The promoters of CASE did not recognize the need for this knowledge. They believed that mechanistic knowledge suffices for developing applications; and, since mechanistic knowledge can be embodied in software devices, they believed that the contribution made by people can be reduced to the knowledge required to operate these devices.

The main purpose of this argument, you will recall, is not to show the absurdity of CASE, but to show how the failure of CASE demonstrates the fallaciousness of all methodologies and definitions – which, in turn, demonstrates the fallaciousness of all mechanistic software theories. For, it is software mechanism – the belief that applications consist of independent structures, which can be fully and precisely specified – that is the fundamental delusion. This delusion leads to the delusion that programming expertise can be replaced with rules and methods, which then leads to the notion of methodologies and definitions, and eventually to CASE. The CASE systems merely implemented

formally what the theories had been claiming all along, what practitioners had been trying before to do manually. So the only logical explanation for the failure of CASE is that these theories are invalid.

The Spread of Software Mechanism

1

The study of software mechanism reveals some disturbing trends. We note a marked shift in our preoccupations: from solving real problems, to dealing with the problems created by the tools we invent to solve these problems. We have to address more and more issues, and these issues are becoming more and more complicated, while moving further and further away from our real concerns. Each issue begets new kinds of problems, which did not exist before but which are now urgent, because they must be solved before we can return to the original problem.

Simple programming problems – which require only a human mind and experience – gave rise to theories of programming, which evolved into complicated methodologies and programming tools, and finally became the monstrous development environments we see today. What started simply as programming became a preoccupation with programming theories and languages, with tools and environments. We also note a tendency to spread this inefficiency into broader domains: from the world of programming itself to the organizations that use software, and then to the rest of society. Issues that ought to concern only programmers end up affecting the users of their applications, who must now deal with some new, software-related problems in addition to their own problems; and these new problems affect then those people with whom the software users are involved. In the end, the entire society is spending more and more of its resources solving spurious, software-related problems, instead of addressing real concerns.

These three trends – making activities more complicated than necessary, shifting the preoccupation from real to spurious issues, and spreading the inefficiency into broader domains – are related, of course. They are different aspects of the same phenomenon, different manifestations of the delusion of software mechanism. Let us briefly examine this phenomenon.

Since the mechanistic software theories are invalid, the software practitioners who follow them keep failing. Their belief in software mechanism acts then as defence, as a way to deny reality. The reality, we saw, is that their failures are

due to incompetence – an incompetence fostered by the very theories they follow. For, in addition to leaving unsolved the programming problems they promise to solve, the mechanistic theories prevent practitioners from gaining the expertise that *could* solve them.

There is no limit to the number of trivial issues that we can contrive if we believe that any problem can be solved by breaking it down into simpler ones. By replacing the original, complex problem with a multitude of isolated and relatively simple issues, ignorant practitioners can delude themselves that they are working toward the solution of the complex problem. They have redefined, in effect, the challenge of programming: from developing software to solve real problems, and at the same time improving their skills, to searching for ways to reduce these problems to simple ones that match their limited capabilities.

The mechanistic doctrine, thus, helps the software practitioners to rationalize their failures. Instead of the difficult task of creating and maintaining applications, they can now perform activities so simple that they cannot possibly fail. The problems that the applications were supposed to address may remain unsolved, but the new problems always have solutions, so the practitioners are pleased with their accomplishments. In the end, the belief in software mechanism permits them to shift the definition of their profession, from solving real problems to solving their own, simpler problems.

What is left is to persuade those who trust them – their employers, and the rest of society – that what they are doing is the utmost that can be accomplished in the domain of programming. But this is easy in a society dominated by mechanistic beliefs. The concepts of software engineering are deemed “scientific” – because based on mechanistic principles – and are therefore readily accepted by everyone. Given our mechanistic culture, it is the *truth* that is hard to accept. We have enough evidence that the most important part of programming is knowledge, talent, skills, and personal experience; in other words, the traditional form of expertise. And yet, it is the fallacious theories and the programming substitutes based on them that are promoted in computer publications, taught by professors and gurus, and endorsed by institutes and associations.

Permitting the software practitioners to fool us with their mechanistic theories is no different from permitting researchers to fool us with mechanistic theories in the human sciences. We saw in previous chapters how the mechanistic doctrine has corrupted academic research in fields like sociology, psychology, and linguistics, turning in effect these disciplines into pseudo-sciences. Invalid mechanistic theories are being pursued for years and decades, and are then abandoned only to be replaced with other mechanistic theories. The “research” performed by these scientists is merely a preoccupation with the spurious, mechanistic problems they themselves keep creating when

breaking down the complex phenomena of mind and society into simpler phenomena. They never explain the actual phenomena; but, despite their delusions and their failures, we continue to trust them, and the universities, and the mechanistic philosophy. It is not surprising, therefore, that we have come to trust in the same way the software practitioners, even as we see them preoccupied largely with the spurious, mechanistic problems they themselves keep creating, and despite their failures.

So those who depend on software learned to deny the reality of software failures just like the practitioners themselves. Software users can see the failure of their applications as clearly as the software practitioners can see the failure of their theories. But, just as the belief in software mechanism helps the practitioners to rationalize their programming failures, this belief helps now those who depend on software to rationalize the failure of their applications.

The software practitioners, we saw, rationalize their failures by expanding the original problem into a vast and complicated array of trivial activities. The real problem – creating and maintaining software applications to address business or social issues – is a well-defined one. It is a difficult problem, however, requiring much expertise, so the software practitioners hide their incompetence by replacing it with a great number of isolated, simple problems. They can then concentrate on *these* problems, and delude themselves that they are making progress toward the solution of the original one. This technique has been so successful in the domain of programming that the *users* of software have now adopted it, in order to rationalize their own failures. Instead of addressing real issues, they are modifying their activities and lowering their expectations to match the inferior applications they depend on.

The software propaganda has succeeded in convincing us to accept a state of affairs that no one would tolerate in other domains. Failures that would be preposterous in manufacturing, or construction, or utilities are considered normal for software applications. The users of applications notice that programmers are using complicated methodologies and development tools, and are calling themselves “software engineers,” and conclude that they are like other professionals. Instead of being outraged by the inefficiency and incompetence they note in software development, the users of applications allowed the creators of those applications to persuade them that software-related problems must form an important part of *their own* preoccupations. Thus, assured that the programming they see is the only kind of programming possible, software users started to expand their own problems into a vast array of trivial, software-related activities. Solving software-related problems, they are now convinced, is the only way to make progress toward the solution of their real problems.

2

When studying the spread of software mechanism, and the resulting incompetence and corruption, it is useful to distinguish three stages. The first stage, now complete, involves the world of programming itself: the programming profession has been almost totally destroyed by the mechanistic dogma, and the inefficiency in those activities directly related to programming exceeds 90 percent (and is sometimes as high as 99 percent). To put this differently, if we had true professionals instead of the present programmers, society would pay less than one tenth of the current cost to derive the same benefits from software as it does now. (In fact, with true professionals we would derive *greater* benefits at one tenth the cost, because the present programmers are far from delivering the best possible applications.)

Programming expertise has been redefined, as I have already remarked, from knowledge of programming to knowledge of ways to *avoid* programming. Solving a problem simply through programming is considered old-fashioned, while attempting to solve it by means of development environments, ready-made pieces of software, and other high-level concepts is seen as modern and professional, no matter how inefficient and expensive is the result. Programming expertise, in other words, means expertise in the use of substitutes for expertise. The responsibility of programmers is limited to operating software devices, or following methodologies, or being aware of the latest theories. They are not accountable for their work: what matters is not whether their applications are useful or not, but whether their activities conform to the current software ideology.

When an application fails, no one is blamed. People cannot be blamed when their responsibility is limited to operating devices, or to following rules and methods. The conclusion is typically that they did not use the latest “solutions” or the latest “technologies.” So the project is simply abandoned, and another one is started – perhaps with a different development environment and different hardware – while all the people involved continue to be trusted and respected. (We will return to this subject in “Software Irresponsibility” in chapter 8.)

A nice demonstration of the first stage is provided by the evolution of Microsoft Corporation – the most aggressive of the software elites. Microsoft started with simple and useful programming tools, but the more successful it became, the larger and more complicated became its systems. Eventually, every tool grew into a huge environment, where the actual programming issues are hidden in a maze of unnecessary features, options, rules, and standards –

all a result of raising the level of abstraction in software development. The effort involved in using Microsoft development environments (and any other environment that depends on the Microsoft operating systems) is due largely to the environments themselves, not our real programming needs. Microsoft has attained, thus, the goal of the first stage: to destroy the possibility of intelligent and responsible programming by forcing programmers to waste their time with spurious problems; to prevent programmers from improving their skills by making their work dependent, not on personal knowledge, but on development environments; and, ultimately, to reduce all programmers, regardless of experience or potential, to mere operators of software devices.



The second stage, now well under way, involves the world of business: corporations, governments, financial institutions, and the like – those organizations that have been heavy users of computers for many years, as well as those that started more recently. Workers in these organizations are undergoing now, in their own fields, the same indoctrination as the one that led to the destruction of the programming profession. And as a result, the same incompetence, inefficiency, and irresponsibility that characterize the world of programming are increasingly affecting all business. Although we are observing now office workers rather than programmers, and the *use* of software rather than its development, the similarity of this stage to the earlier one is striking.

The second stage started when the users of software realized that programmers did not, in fact, provide an acceptable level of service: new applications took too long to create, or were never finished, or were inadequate; simple modifications that should have been implemented in a day or two took months, or were never done, or did not work. Frustrated by this state of affairs, and convinced that this was the only kind of programming possible, the users fell into the same trap as the programmers themselves: they accepted the solutions proposed by the software elites – the same solutions that had led to programming incompetence, and hence to their current problems, in the first place. As before, the solutions consisted, not in improving anyone's skills, but in means to avoid the need for programming. Incredibly, the elites turned now to the users themselves, offering them software devices that would eliminate their dependence on programmers: ready-made applications and, when these proved to be inadequate, user-oriented software development tools.

Asking users to depend on generic applications, or to develop their own, was tantamount to acknowledging the incompetence of programmers, and hence the failure of the mechanistic programming theories. No one recognized this obvious fact, however. So the same elites that were preventing

programming expertise were seen now by their victims – by the users of software – as saviours. The incompetence of programmers, we saw, is due to the mechanistic theories, which force them to depend on aids and substitutes instead of improving their programming skills. They are restricted to mechanistic knowledge, which is the same as staying forever at the level of novices. By invoking the same theories, the elites could now deceive and exploit the *users* of software, and prevent expertise also in other occupations.

The first end-user aids were simple software tools for database query, reporting, spreadsheets, and the like. These tools, needless to say, could not fulfil the promises made for them: they could not be a substitute for serious applications. So the users found themselves dependent on these aids, and on the software companies behind them, even as their needs remained unsatisfied. But, again, no one questioned the software concepts leading to these aids, nor the honesty of the elites. It was the aids that were suspected, so the same process of expansion that had occurred earlier for the programming aids started for the end-user aids. The simple tools grew into large office systems, complete with programming languages, heavy instruction manuals, and proficiency courses. Each new version made them more complicated, until they finally became the monstrous software environments we see today.

These systems, clearly, are the counterpart of the development environments employed by programmers. Also like the development environments, they are fraudulent, in that they do not eliminate the need for programming expertise: to create serious applications, their users need almost the same knowledge as they would if relying on programming languages. They are called *productivity* systems, but this is a spurious term, as its true meaning is the opposite of its literal one. The purpose of these systems is to *reduce* productivity, by consuming everyone's time with software-related activities. And they accomplish this by replacing the dependence on personal knowledge with a dependence on software devices. Like programmers before them, workers in various fields are now prevented from practising their profession and from improving their skills. Increasingly, the only thing they learn is how to operate software devices: how to select and combine functions from a range of alternatives.

Just as the programming environments can replace only the *easy* parts of software development, the office systems can replace only the *easy* parts of software use – the parts requiring *mechanistic* knowledge. No matter how elaborate they are, these systems cannot replace the parts demanding *non-mechanistic* knowledge, and hence human minds. The real problems, thus, remain unsolved, but everyone believes that the only answer is to acquire even more software devices. Like programmers, office workers now remain at novice levels, because, instead of practising, they waste their time with the

problems generated by their software tools. Just like programming expertise, the expertise of office workers has been redefined to mean the skill of using substitutes for expertise.

In one occupation after another, the responsibility of people is being lowered, from the traditional one – solving a real problem, making a valuable contribution – to merely knowing how to deal with software-related issues. In some occupations, the simple skills needed to operate a particular software device – skills that can be acquired by almost anyone in a few days – are already considered more important than a lifetime of experience.

A demonstration of the second stage is provided, again, by the evolution of Microsoft Corporation. Having destroyed the concept of expert and responsible programming, Microsoft turned to the world of business. The purpose of its office systems is to raise the level of abstraction in office work; to enforce the perception that any task can be reduced to a combination of some built-in functions, thus making everyone dependent on software devices; to waste workers' time with the problems generated by this dependence, thereby preventing them from gaining knowledge and experience; and, ultimately, to reduce all workers, regardless of position or skills, to mere operators of software devices.



The third stage involves the spread of software mechanism into the rest of society – into our homes and personal affairs, in particular. This stage has just begun, and it is hard to predict how software mechanism will affect our life; that is, to predict what will be the equivalent, in our personal affairs, of the incompetence, the inefficiency, and the corruption that we see now in programming and in business. Judging by the previous stages, however, we can expect an ever-growing dependence on software environments and on the companies behind them; specifically, a dependence on systems that promise to solve our important problems while addressing in reality only the simple ones. In the guise of information, or education, or entertainment, more and more types of software devices will be invented, desktop and mobile, and we will spend more and more of our time with the trivial preoccupations engendered by their use.

Thus, we will have fewer and fewer opportunities to use or to develop non-mechanistic knowledge. Our definition of expertise, creativity, and responsibility will be degraded, in everything we do, to mean simply the skill of selecting and combining the functions provided by software devices. The goal of the elites now is to make all human activities as insignificant and inefficient as they have made programming and office work; to make us all as ignorant as

they have made programmers and office workers, by forcing us to depend on their systems instead of improving our minds; and, ultimately, to reduce all human beings on earth to mere operators of software devices.



The three stages overlap, of course, and this analysis may not be entirely accurate. The distinction is still useful, however, because it helps us to observe the progression of software mechanism: its expansion, outward, from the narrow domain of programming into all aspects of human existence. By studying the first stage, which is now complete, we can perhaps foresee the evolution of the other two stages. (This study is the subject of the next chapter, “Software Engineering.”)

We should try to imagine a society where all people are as inefficient and irresponsible in their pursuits as the software practitioners are in theirs today. And, by studying the past and current rate of expansion, we should try to estimate how long it will take the software elites to bring about this condition. The frightening conclusion would be that, within a few decades, the second and third stages will also be complete. So we should ask ourselves whether a modern society can function at all when everyone’s knowledge is at the low levels we see today in programming. Programmers can be as incompetent as they are only because the rest of us are willing to pay the cost of their inefficiency; that is, willing to support a software bureaucracy. But if we, too, were to be like that, who would support us all?

It is significant that the software elites are not content with merely *exploiting* society – something they have already accomplished, in the first stage. Their objective is not just to extract, through the software bureaucracy, vast amounts of money from society. In their plan of domination, the second and third stages are as important as the first, and they will not stop until every person is turned into an active member of the software movement.

Like all totalitarian ideologies, the software revolution is a mass movement. It is not founded on authoritarianism, but on mind control: people are not threatened, but indoctrinated; they must become devoted followers, and participate, on their own accord. Physical force is needed only against those who, despite the indoctrination, still fail to appreciate the benefits of the new social order.

Viewed in this light, the three stages of the software movement parallel the evolution of political totalitarian movements like Nazism and Communism. The revolution starts with a core of believers – who eventually become the elite, or the Party – and spreads outward, bringing increasingly broad segments of the population into its ranks. The revolution cannot end until every person and

every event in society conforms to its ideology. But because the ideology is based on pseudoscientific, fallacious notions, it cannot actually work. The initial phase, when only a small portion of society is involved, usually appears successful, and this gives the believers confidence in their utopian visions. (The software movement is currently in this stage.) The initial phase appears successful, not because the ideology is valid, but because the few who embrace it deceive and exploit the rest of society. As the movement spreads, however, more and more people are turned into bureaucrats who merely serve the ideology, and fewer and fewer are left who do real work and can be exploited. Society becomes increasingly corrupt and inefficient, and eventually destroys itself. (See also the related discussion in the introductory chapter, pp. 30–31.)

