

Theory of Computing: A Scientific Perspective¹

(preliminary version)

Oded Goldreich²

Avi Wigderson³

October 19, 1996

¹Available from URL <http://theory.lcs.mit.edu/~oded/toc-sp.html>

²Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, ISRAEL. E-mail: oded@wisdom.weizmann.ac.il

³Institute for Computer Science, Hebrew University, Givat Ram, Jerusalem, ISRAEL. E-mail: avi@cs.huji.ac.il

Abstract

We are gravely concerned with the contents, spirit and recommendations in the Aho *et. al.* Report on the Theory of Computing (TOC). This report “assesses the current goals and directions of the Theory of Computing (TOC) community, and suggests actions and initiatives to enhance the community’s impact and productivity”. However, we believe this assessment to be inaccurate and the main recommendations to be wrong and extremely dangerous to the very existence of TOC.

In this essay we provide an alternative assessment of TOC, namely as a fundamental scientific discipline rather than as an engineering subcontractor. We argue that the TOC great achievements, productivity and impact so far (both scientific and technological) were due to the autonomy it had to pursue its intrinsic goals. Our main recommendation is that **in order for TOC to prosper in the coming years, it is essential that Theoretical Computer Scientists concentrate their research efforts in Theory of Computing and that they enjoy the freedom to do so.**

We base our analysis on clearly stated (and generally accepted) philosophical beliefs, and concrete evidence from the history of our field. We provide a critique of the Aho *et. al.* Report from this basis. Another important part of our essay is an attempt to explain (and thereby reduce) the sense of frustration of some members in our community with respect to the status of basic research in TOC.

Our hope is that this essay will stir up a serious discussion in the TOC community, focused on science. We further hope that such free discussion will result in much stronger self-esteem and belief in the important scientific role of our field. Armed with these, the TOC community will be able to do much better research and education as well as deal with (externally imposed) conditions regarding funding and jobs.

Contents

1	Introduction	2
2	On the merits of TOC	5
2.1	Culture, Science and Technology	5
2.2	Evaluating (the importance and success of) scientific disciplines	6
2.3	On the fundamental nature of TOC and its success so far	6
2.4	On the impact of TOC on Technology	9
2.5	On the impact of TOC on other sciences	11
2.6	On the future of TOC	12
3	The true problems of TOC	13
3.1	Deep unjustified frustration	13
3.2	Lack of a leadership group	15
3.3	The outside pressures	16
4	Critical reading of the Aho et. al. Report	17
4.1	Authority and representation	17
4.2	The relation of TOC to CS and other disciplines	18
4.3	Future Prosperity	19
4.4	Communicating the fundamental nature of TOC	20
4.5	Evaluating the achievements of TOC	20
4.6	Evaluating the technological impact of TOC	21
4.7	The value structure of TOC	22
4.8	History according to the Aho et. al. Report	23
4.9	TOC according to the Aho et. al. Report	23
4.10	Summary	25
5	Our Recommendations	26

Chapter 1

Introduction

We are gravely concerned with the contents, spirit and recommendations in the Aho *et. al.* Report on the Theory of Computing (TOC). In this report, some very prominent Theoretical Computer Scientists, despite their good intentions and in the face of external pressure and internal frustration, sign what may become a death (by suffocation) warrant to their community. Internally, it calls upon the best creative forces in the TOC community to direct their efforts away from its major scientific goals and into application areas¹. Externally, it legitimizes funding agencies, deans and department chairs to reprimand those who fail to follow this direction². All this is taking place at a time when TOC has been constantly producing exciting, fundamental research of the greatest importance (to Computer Science as well as other disciplines), but is still far from achieving its scientific goals and dearly needs the best creative forces within it. In these circumstances, such a report is unprecedented (to the best of our knowledge) in the history of science.

The point of departure of the Aho *et. al.* Report is the current funding and job situation, and its assessment of TOC is focused on the impact of TOC on computer technology and other disciplines. In this essay we offer an alternative view for assessing TOC, namely the scientific one, and our evaluation and recommendations drastically differs from the main message of the Aho *et. al.* Report [1]. True to our nature as theoreticians, we explicitly recall the philosophical foundations on which any assessment of a scientific discipline should be based. Using these criteria, we evaluate the importance of TOC as well as its achievements so far. We also provide a critique of Aho *et. al.* Report [1] which is based on the same foundations.

Our fundamental theses are:

- TOC is the science of computation. It seeks to understand computational phenomena, be it natural, man made or imaginative. TOC is an independent scientific discipline of fundamental importance. Its intrinsic goals (those which were achieved, those which are yet to be achieved, and those which are yet to be defined) transcend the immediate applicability to engineering and technology.
- Research in TOC has been extremely successful and productive in the few decades of its existence, with continuously growing momentum. This research has revolutionized the understanding of computation and has deep scientific and philosophical consequences which will be *further* recognized in the future. Moreover, this research and its dissemination through education and interaction was responsible for enormous technological progress.

¹Indeed, some have already confessed to us that they feel forced to do so.

²Indeed, some had already gleefully picked up this hammer and started applying it far beyond the naive intentions of the authors of [1].

- The success of TOC is directly correlated to the extremely high quality and creativity of researchers in TOC, to their independence, and to the fundamental (and exciting) nature of the questions TOC addresses.
- In order for the Theory of Computing to prosper in the future it is essential that TOC attracts the same calibre of researchers, that Theoretical Computer Scientists concentrate their research efforts in Theory of Computing and that they enjoy the freedom to do so.

It should be clear that free pursuit of their research interests may well lead individual scientists to work closely with/in application areas. Indeed, our field has already an admirable tradition where many TOC leaders (undirected and unforced) chose to redirect their research so as to strongly influence application areas. Yet, decisions taken by individual scientists following their own understanding of the discipline differ drastically from attempts to direct the whole discipline towards directions which are not intrinsic to it. Thus, we **completely reject** the main conclusion in Aho *et. al.* Report [1] that the prosperity of TOC *depends* on service to other disciplines and immediate applicability to the current technological development.

We feel that the wrong conclusions of the Aho *et. al.* Report [1] follow naturally, not only from external pressures, but from what we view as misconceptions and dangerous moods that have spread through the community in the last few years, which the report sadly echoes and amplifies. We feel that it is extremely important not only to criticize these misconceptions and moods, but also to try to explain their origins. We trace these origins to two sources. The first source is a deep (but unjustified) feeling of frustration among the members of the TOC community. The frustration is due to unrealistic expectations by which the TOC community should have been able to gain by now an almost full understanding of the nature of efficient computation. The unrealistic nature of these expectations stems from the unimagined (by the founders of the field) depth and richness of TOC, revealed in the last 20 years. The second source is the lack of a “leadership group”, deeply convinced of the importance of the discipline and continued research in it, which is willing and able to oppose pressures from the outside, as well as further this conviction to the junior TOC generations.

To summarize, as the Aho *et. al.* Report [1] says, “trouble is looming for the field of TOC”, but ironically it stems more from the report itself and the moods it reflects than from the outside pressures. We call all of you to examine the main theses of this essay. We believe that a deep conviction in these theses is essential for successfully confronting the dangerous internal moods as well as the outside pressures.

The primary target of our essay is the junior scientists in TOC. We fear that they have been greatly demoralized and damaged in the last years and we believe that they can recover if they realize the importance of their historical role. We call upon each of them to determine his/her goals and directions in TOC based on his/her own understanding of TOC. We direct this essay also to the senior scientists in TOC, with the belief and hope that they will support any good work in TOC as long as it is governed by a candid wish to advance TOC.

Organization: In Chapter 2 we justify our conviction that TOC has a major role and place among the sciences. In Chapter 3 we try to explain the reasons for a feeling of crisis within TOC and offer ways to overcome these feelings and confront outside induced problems. In Chapter 4 we provide a critique of [1]. In Chapter 5 we present our recommendations.

Important comment: After reading the Aho *et. al.* Report [1] a month ago, we have discussed it with many colleagues. All have found major flaws and dangers in it, but none seem to plan a

significant rebuttle. This state of affairs, compelled us to write this essay quickly and in a sense of urgency so that it is ready by *STOC96*. Despite our impression that the views expressed in this essay enjoy great support in the community, we want to make it clear that we alone take full responsibility for its contents. The time pressure, as well as our personal biases and limitations, make this essay far from perfect, even in our eyes. Nevertheless, we hope it will suffice to stir up a true open debate in the community, in which all silent voices will become vocal. We plan to invest more time into improving this essay, and integrate comments and criticism of others into it.

Literary comments: In this essay we have made several “literary” choices that may harm its effectivity. One was a sparse use of satirical tones – this was out of despair. Another is the use of examples – these are always easier to attack than the fundamental claims, but were nevertheless provided as they tend to illustrate and facilitate abstract arguments. A third is repetition of our basic claims – while boring, this was a reaction to their (near) absence in Aho *et. al.* Report [1].

A personal comment: We consider ourselves very fortunate to have taken our first steps in the Theory of Computing in an enlightened and exciting atmosphere, very much different from the current mood. We consider this essay as a minor payment towards our duty to try and provide a similar atmosphere for the new generations of TOC students.

Chapter 2

On the merits of TOC

We start by explicitly stating our beliefs regarding science and the evaluation of scientific disciplines. These beliefs are far from being original. They are rooted in the best philosophical and scientific traditions. For lack of time and energy, we do not provide a host of references.

2.1 Culture, Science and Technology

The search for truth and beauty is the essence of civilization. Since the Renaissance, the search for truth takes the form of (or is called) Science. Technology is an important by-product of the scientific progress, not its *raison d'être*. Furthermore, philosophical reasoning as well as experience show that technology is best served by a free scientific process; that is, a scientific process which evolves according to its own intrinsic logic and is not harnessed to the immediate technological needs. Such free scientific process evolves by formulating and addressing intermediate goals which are aimed at narrowing the gap between the ultimate goals of the discipline and the understanding achieved so far.

It is ironic that as the contribution of science to technology becomes wide-spread, a popular demand arises to have more. Namely, the success of science and in particular the benefits of its technological by-products causes the populace to turn against science (in the form of demands that science deliver even more consumable commodities). Still, one has to oppose these demands. Science is to maintain its *autonomy* which is correlated to its success. In the long run, this is also the best way to serve technology.

Technology evolves mostly via applied scientists and engineers who use the scientific knowledge they have acquired and their own creative forces to the development of specific applications. Contrary to popular beliefs, the most important contributions of science to technology do not stem from the harnessing of scientists to engineering tasks, but rather from the fact that scientists instruct and enrich the thinking of these engineers. The education of engineers does not reduce to the acquisition of information. Its more important features are the development of conceptualization and problem-solving abilities. The conceptual frameworks of the discipline are offered to the student and the better these frameworks are the better an engineer he/she may become. This form of education is most effective when done by good scientists who enjoy the freedom to pursue their own research interests.

It is important to note that the nature of the process by which science effects technology makes it very hard for the laymen, and sometimes even the expert, to trace a technological breakthrough to its scientific origins. Almost always these breakthroughs depend on the conceptual scientific

framework and very often they utilize specific discoveries which were considered totally impractical at the time of discovery (e.g., complex numbers and electricity).

2.2 Evaluating (the importance and success of) scientific disciplines

The scientific disciplines are defined by the questions they address. The *importance of a discipline* is determined by the nature of its formative questions. The more fundamental these questions are the more important the discipline is. Educated laymen and certainly scientists can usually assess how fundamental major scientific questions are.

The *success of a discipline* is measured by the progress it achieves on its own formative questions. To measure the amount of progress one has to understand the questions and the state of knowledge of the discipline with respect to these questions. This usually requires the understanding of experts, but can be conveyed to scientists of other disciplines.

Neither the importance nor the success of a scientific discipline can be measured by the impact of its current discoveries on technology (or on other disciplines). If the discipline is indeed important and successful such impacts are likely to follow. However, rarely will this impact be linearly related to the scientific progress in the discipline.

Individual scientific disciplines do not exist in a vacuum. The healthy evolution of a scientific discipline is sensitive to *scientifically relevant* inputs from other disciplines as well as technological developments. We wish to stress that the influence of these inputs is determined by the disciplines internal logic and inherent goals and that such influences are vastly different from non-inherent suggestions (e.g., that in order to increase funding and/or employment opportunities the discipline should pursue alternative directions).

2.3 On the fundamental nature of TOC and its success so far

The Nature of Efficient Computation and its natural as well as surprising derivatives, is the formative question of the Theory of Computing (TOC). We consider this question to be one of the most fundamental scientific questions ever asked. Unfortunately, the fundamental status of this question is usually disregarded due to its immediate technological impact.

We feel that both the fundamental nature of the questions of the Theory of Computing and the success of our community in engaging these questions (up to this very day) are evident. To be on the safe side, here is some evidence.

An excellent demonstration of the the fundamental nature of TOC was provided by Papadimitriou [2] in his survey on the impact of NP-completeness on other sciences. Papadimitriou lists about 20 diverse scientific disciplines which were unsuccessfully struggling with some of their internal questions and came to recognize their intrinsic complexity when realizing that these questions are, in some form, NP-complete. According to his bibliographic search, NP-completeness is mentioned as a keyword in about 6,000 scientific articles per year. How many scientific notions have had such impact?

More generally, TOC has established a direct relationship between structural and computational complexity. Efficient algorithms are discovered almost only if tangible mathematical structure exists. This connection has already benefited mathematical progress in many areas such as Number Theory, Algebra, Group Theory and Combinatorics, where on one hand a need for efficient algorithms existed, and on the other hand the search for them has generated structural results of

independent interest.

Actually, we tend to forget the revolution in problem-solving introduced by the TOC treatment of algorithms. This revolution consists of the explicit introduction of the concept of an algorithm and the measures for its efficiency, the emphasis on data representation and organization, the general techniques for creating algorithms for classes of problems, and the notion of reductions between problems. Needless to mention the impact of all these on computer practice, but we wish to stress the impact on any kind of problem solving.

The TOC has drastically changed the perception of knowledge and information. Specifically, the TOC stresses that different representations of the same information may not be *effectively* equivalent; that is, it may be infeasible to move from one representation to the other (although a transformation does exist). In this new world, publicly available information may be unintelligible. All of Modern Cryptography is based on this Archimedes' point, and its scientific and technological impact are well known. Here we wish to suggest that this revolution applies not only to computer systems but to any aspect of human interaction in which privacy and fault-tolerance are important concerns.

The TOC has introduced totally novel ways of understanding and using randomness. The probabilistic algorithms developed within the TOC use randomness in many varied sophisticated ways. The applicability of randomized procedures for solving tasks from different domains such as number theory, optimization and distributed computing is amazing. Moreover, the growing study of derandomization has led to derivation of better deterministic algorithms from probabilistic ones.

Combining randomness and interaction lead TOC to create and successfully investigate fascinating concepts such as interactive proofs, zero-knowledge proofs and Probabilistically Checkable Proofs (PCP). Each of these concepts introduces a deep and fruitful revolution in the understanding of the notion of a proof, one of the most fundamental notions of civilization. Furthermore, these revolutions bore fruits which reached far beyond the realm of proof systems. For example, work on PCP led to the first breakthrough in the understanding of the hardness of approximation. This is but one incredible demonstration of the how probabilistic thinking leads (very indirectly and non-trivially) to fundamental understanding of totally non-random phenomena.

In addition, combining randomness and complexity, TOC has generated meaningful notions of pseudorandomness. Computational hardness yields pseudorandom generators: using "one-way" functions, randomness can be "stretched" in an almost unlimited way as far as efficient observations are concerned. This yields the stunning (to most scientists) conclusion that if their Monte-Carlo algorithm (estimating perhaps a numerical integral or simulating a physical process) behaved differently on sequences produced by such generator, than on genuine random sequences, then they have discovered an efficient factoring algorithm! Totally different pseudorandom generators which TOC discovered can fool any space limited algorithm. Since all standard statistical tests have such implementations, this is great news to Statisticians, Physicists, and most Social Scientists who use such tests on everyday basis. Namely, the results of all their experiments are guaranteed to hold even if they replace all their random choices by pseudorandom choices produced by from tiny random seed.

TOC has gained considerable understanding of organizing work on huge systems of many components. The study of parallel algorithms resulted in amazing ways to get around "inherently sequential" tasks. Subdividing work to smaller chunks in efficient and balanced ways is taking place not only in computer systems but in many organizations, and the insights gained by TOC are avail to them too. A different kind of parallel computing arises in settings where the information is distributed among the components of the system. TOC studies of such distributed environments resulting in models and methods of consistency, recovery, knowledge, synchrony and decision mak-

ing, are relevant not only to (distributed) computer systems but also to economics and other social sciences.

The organization and availability of information was always a major part of civilization, and in particular science and technology depend on it. The models and solutions developed by TOC for such problems not only resulted in computer systems that would do it for people, but in the very way people and institutions have to think about information. The amazing new abilities to handle huge masses of data increase, rather than decrease, the human decisions on what they want to be stored, what access patterns they want to allow and disallow, what should be retrieved quickly and what can take longer, etc. The theoretical understanding enables to formalize their demands, and enable programmers (who should understand the algorithms and data structures as well) either to satisfy these demands or to explain why they are impossible to achieve.

Likewise, some of TOC's insights to performance analysis, the minimizing and balancing of several resources, are of universal applicability. One example is the notion (and techniques) of competitive analysis, whose applications range from operating systems to information compression to emergency services to stock-market investments. More generally, asymptotic analysis has taught us that structure is often revealed at the limit. The adversarial point of view developed for worst case analysis (both of inputs to algorithms and behavior of distributed systems) has taught us a similar lesson: structure is often revealed under the worse circumstances and may be obscured by unjustified assumptions on "typical behavior". Such structure often leads to better (in every respect) theoretical and practical solutions.

Finally, let us mention that many inter-disciplinary scientific activities involve and further seek the participation of TOC members. These include the different "neurocomputational" groups (encompassing brain models, learning, and neural networks, involving physicists, biologists, psychologists) and "rational behavior" groups (encompassing economy, ecology, evolution, competition, and decision making, involving economists, statisticians, psychologists and mathematicians). They want TOC to be there since they have recognized, in contrast to some members of the TOC community, the universal value of the problems TOC deals with and the understanding TOC has obtained so far, and in particular their relevance to these areas.

Clearly, lack of space, time and knowledge prevents us from going on. Still, the massive list above illustrates the fundamental nature of our endeavours from the scientific point of view. But they are fundamental also from two other important viewpoints. One is the philosophical viewpoint, which has dealt with many of the notions and questions above for centuries, and which receives a fresh, radically different perspective (namely the computational one) from TOC. As an example consider the question of P vs. NP vs. CoNP. Some tend to think of it as a mere technical question and miss its deep philosophical significance: Understanding the relation between the difficulty of solving a problem to the difficulty of verifying the correctness of the solution, to the difficulty of proving that no solution exists. Additional examples are the TOC perceptions of the notion of a proof, its view of randomness, and its emphasis on the importance of specific representations. The second viewpoint is the potential contribution of TOC to the general education and enrichment of humanity. Many notions, problems and even some of the solutions TOC has produced are available for understanding (in nontrivial levels) by laymen. We have successfully tried to explain some of them to elementary school kids (and indeed we foresee some of them taught and used as teaching paradigms in grade and high school). Few sciences (which existed for many centuries) can compete on these grounds with what TOC achieved in a few decades.

To summarize, this subsection illustrated the fundamental importance of TOC as well as its success. As for the latter point, let us stress that the achievements sketched above are more or less equally spread over the last 30 years, and many are very recent. Indeed, the rate of progress done

by TOC in these years is astonishing and there is no inherent reason for this progress to stop. It is thus essential to oppose the external pressures and internal moods which endanger the continuation of the fundamental and successful research in TOC.

2.4 On the impact of TOC on Technology

While we rejected technological impact as a measure of importance and progress of a scientific discipline, the enormous impact of TOC research on technology should not be made a secret. We are far from experts regarding this impact, still there are a few points that even we can tell. We hope and believe that a much better treatment will be given in the future by more qualified colleagues.

The most important impact of TOC on Computer Science and Technology stems from the fundamental goals of TOC. In its endeavour to understand the nature of computation, TOC created general abilities to conceptualize, model, unify, solve and analyze computational mediums and problems. The effects of this understanding are present in essentially every working system and algorithm on earth. Without them the computer revolution, which has changed life on this planet in a fundamental way and will continue to effect it at increasing speed, would simply not be possible! Indeed, they are the very reason that theory courses are mandatory for all undergraduates in computer science departments. They are the reason that most applied computer science courses are not a mere collection of ad-hoc tricks and are thus suitable to be taught in universities. They are the reason that the originators of technological breakthroughs, as well as all engineers and programmers, can actually think, talk, present and evaluate their ideas. Some critics may say that these understandings were achieved long ago, and there is no need for further “refinements”. This is contradicted by many technological advances which have resulted (and will continue to result) from *recent* developments of such understandings regarding, for example, parallel, distributed, interactive, secure and fault-tolerant computation. Many such developments were achieved by special interest groups within TOC, who took on to study in depth such models and algorithms. Their specialized conferences, which are a relatively recent phenomena, often enjoy the active participation of more applied scientists, who have both easy access to this knowledge as well as a forum to influence its direction.

It is crucial to recognize and communicate the fact that most of this understanding resulted from attempts to solve a concrete problem under particular technological constraints. Rather, it came from generalizing the problems and abstracting away unnecessary technological details to the point that enables finding structures and connections to other knowledge. Only then could applied scientists and engineers, who had both the theoretical understanding as well as the mastership of the specifics of the technological task, fuse them together to a successful practical object. The value of this approach has many examples, and we discuss only one.

- **NC and the PRAM model.** As an example, we choose on purpose the class NC and the PRAM model, a common bashing target of “practical people” (as well as of the TOC self-destructive fashions). While the direct applicability of this model (and algorithms for it) may still be controversial, several parallel systems builders we have talked to have totally changed their attitude towards it. Technological changes have made it much closer to reality than realized 15 years ago, which teaches us a moral regarding fine tuning of theoretical models to current technology in a field in which the latter is changing at such rapid speed. But this retrospective fact is not the source of the value of PRAM. Its value stems from the fact that it is a good framework for developing an understanding of the paradigm of parallel computing. Specifically, the answer to the bashers should have always been that a very fast

practical parallel algorithm (for, say, sorting or matching), on a particular architecture (say Connection Machine), is almost necessarily also an NC algorithm on PRAM. If we cannot find the second, how can we develop the first. Moreover, while a PRAM algorithm may never be implemented as is, the algorithmic techniques, communication paradigms and data structures developed in its study, have strongly influenced many different practical systems.

In general, one should advocate the value of abstractions which address some fundamental aspects of an important problem (even if they seem not address all aspects), and warn against the shortsightedness captured by dismissing such abstractions as irrelevant. The study of such an abstraction is more likely to yield fundamental insights than the study of the “real problem” (assuming such a creature exists – actually there is never one real problem but rather many different related real problems and what these have in common may well be the dismissed abstraction). Only later will people, with a concrete application and technology in mind, be able to fine-tune the theoretical understanding to their needs. (This in itself may require significant research and implementation, that was and is taking place by computer scientists and engineers, and which resulted in so many successful technological developments.)

It is equally important to recognize and communicate that it was the freedom and time given to TOC researchers to pursue these general directions, in real attempt to understand novel computational media, that resulted in such progress – quite often in surprising and unexpected ways.

One can illustrate the point above by numerous examples, and in particular the ones given in Aho *et. al.* Report [1] (regrettably without such illustration) are perfect. We prefer to give two very recent examples whose technological and practical effects are imminent and yet to come. So far their “practicality” is demonstrated by a major leap in the algorithmic understanding of major problems. This leap is rooted in developments of complexity theory which, at first and for a long time, seemed totally irrelevant to the latter or any other algorithmic task. Such leaps are frequent in our field, and are due to the freedom of pursuing scientific intuition, as well as to the strong communication and information exchange between the various subareas of our field.

- **The Euclidean TSP Algorithm.** A few weeks ago Sanjeev Arora announced a polynomial time approximation scheme for the Traveling Salesman Problem (and a host of other combinatorial optimization problems) in the plane. The problem itself was a major object of study in our field for decades. The failed attempts to find such approximation scheme resulted in fundamental contributions to NP-completeness, probabilistic analysis, approximation algorithms and mathematical programming. It also resulted in enormous efforts to understand the relative power of various heuristics.

The techniques present in the algorithm of Arora were available decades ago! Why was it only found now? While this is a source of speculations, Arora himself tells how he came about it. The algorithm arose from his attempts to generalize the inapproximability results of metric TSP to Euclidean TSP, attempts which revealed to him the extra structures of the Euclidean case. These attempts were based on the surprising connection of PCP proofs to hardness of approximation. In turn, these “mysterious” proofs arised from abstract results like MIP=NEXP (relating “clearly impractical” complexity classes). Moreover, the MIP model of multi-prover interactive proofs was suggested by Shafi Goldwasser as a generalization of interactive proofs (themselves the outcome of amazing developments). Needless to say that Goldwasser did not think of approximation algorithms when she suggested the new model.

- **Efficient Error Correction.** It was only a year ago that Dan Spielman discovered a *linear-rate* code which has asymptotically optimal (i.e., linear time) encoding and decoding algorithms. This central problem of communication, that originated with Shannon half a century

ago, has attracted the best minds in Information Theory, Mathematics, Electrical Engineering and Computer Science, and has resulted in beautiful and important theory. Still, this major problem, resolved by Spielman, was beyond reach.

The construction of Spielman closely mimics the construction of a superconcentrator. This object was not available to most scientists working on this problem, and Spielman learned about it from Complexity Theory. The superconcentrator was invented in TOC, by Valiant, in his attempts at one of the quintessential impractical problems – proving circuit lower bounds. Failing to do that, Valiant turned to an even more impractical problem – to show that this particular attempts will necessarily fail! Here he was successful. He (noconstructively) exhibited the existence of expanders, and used them as building blocks of linear size superconcentrators. A deep and beautiful mathematical theory developed, motivated by the explicit and efficient construction of expanders, which effected diverse areas of TOC. More to the point of this subsection, indirectly and through much further work, derivatives of the study of expanders became extremely relevant to technological development concerning communication networks and protocols for a variety of parallel and distributed architectures.

It is our opinion that the amazing scientific consequences and the surprising practical implications which sprouted (and will continue to grow) from the totally abstract and impractical proposals of Goldwasser and Valiant in the examples above, are *alone* well worth the meager investment so far of the world in TOC. It is sad to speculate that in the world where the Aho *et. al.* Report [1] is adopted (as one of its author conceded), such proposals are likely to be rejected on the basis of “plac[ing] excessive weight on mathematical depth and elegance, and attach[ing] too little importance to genuine links between theory and concrete applications”, especially if written by junior people.

2.5 On the impact of TOC on other sciences

In the short time of its existence, TOC has had an unprecedented effect on other sciences. This has taken at least three forms.

- **Algorithms.** Many sciences use heavy computation for their research, mainly for simulation and analysis. The advances in fundamental algorithms in TOC, on data structures and general techniques are essential for them to understand, so as to optimize their computational resources. The impact of these on the rate of progress in these sciences cannot be underestimated. Moreover, sometimes such disciplines generate a particular type of problems for which the general algorithmic knowledge does not suffice. In some cases where these problems raised sufficient scientific interest (perhaps luckily timed with internal developments), TOC was quick to pick up and study its natural computational structure. Two such superb examples are the great advances TOC has made in understanding and analyzing random walks, so often at the base of simulations in Physics, and its contributions to number theoretic and algebraic algorithms.
- **Natural Computational Models.** Nature computes! While this was observed long before computer science existed, TOC supplied the mechanisms to model, discuss and explain these phenomena. A recent challenge directed by TOC towards Physics is whether a Quantum Computer can be built? But even without the demonstration of the excessive power of the Quantum Computer model (e.g., Shor’s polynomial-time Quantum algorithm for factoring),

we speculate that complexity may be the right way of thinking about decoherence of a quantum mechanical system. The brain is another computational device whose understanding seems to be extremely far, but to which our unique contributions in neural networks and computational learning are providing important stimulation. Another natural source of (biological) computation, based on progress in molecular biology, was discovered by TOC and is studied with at least some interesting potential.

- **Universality of TOC notions.** As pointed out in Section 2.3, the unique computational point of view of TOC and its conceptual derivatives, has resulted in surprising impact on intrinsic studies of other disciplines. NP completeness, discovered over 20 years ago, has had a sweeping effect. But our view on other notions such as randomness, pseudorandomness, interaction and approximation is only beginning to take effect.

It should be reiterated that the discoveries above has made a fundamental impact on these sciences, and have lead them to reassess their points of view on some basic intrinsic questions and pursue novel research directions. We wish to stress that, having sound tradition and self esteem, these sciences were not (and could not have been) forced to pursue these novel directions by TOC or anyone else. Their choice was based on their scientific understanding of their intrinsic goals. Similarly, the interest of TOC in these problems arose from the understanding of TOC researchers that these problems are relevance to the goal of understanding computation. The amazing success of this impact and the high and growing regard to TOC in these sciences, again, stems from the intellectual freedom in which these interactions arose. Again, even a small fraction of these effects justified the investment so far in TOC.

2.6 On the future of TOC

We believe that the notion of efficient computation will further revolutionize the way people think about problems and in particular the way scientists think about basic problems in their disciplines. It is hard to imagine the effect that a *deep* understanding of efficient computation may have on the thoughts of people in the future. To have even more impact on the sciences, TOC has to get a better understanding of the nature of efficient computation, and the other sciences have to further discover the relevance of these notions and results. When this will happen, these sciences will seek insight to computation and if TOC will not commit suicide in the meanwhile it will be there to provide it.

Chapter 3

The true problems of TOC

This section is admittedly speculative, but so are the relevant parts in Aho *et. al.* Report [1]. Needless to say that we believe our speculations and reject alternative speculations offered by the Aho *et. al.* Report (see Chapter 4).

The research community in Theory of Computation (TOC) seems to be a crisis in the last few years. Part of the crisis is induced from the outside: the anti-intellectual atmosphere spreading in the U.S., the ending of the cold-war economy, and the end of the exponential growth of Computer Science departments throughout the world. One can do very little about these factors, except recognize them, adapt expectations accordingly, and devise ways to ensure the momentum of TOC research under less favorable conditions. However, in addition to the outside factors, we distinguish two internal factors specific to TOC. One is a deep (unjustified) feeling of frustration and the other is the lack of a “leadership group” which is willing and able to oppose pressures from the outside. We elaborate on both factors below.

3.1 Deep unjustified frustration

Frustration is an unavoidable feeling which accompanies any challenging human enterprise; specifically, Science and Art. We claim that the TOC community is more frustrated than normal (in other disciplines) and that this frustration is unjustified. We trace the reason for the deep frustration to the gap between the unrealistic expectations of the founding fathers of TOC and the current state of TOC. It seems that many of the founding fathers did not expect TOC to have much depth and did not intend it to become an independent scientific discipline. In contrast, they and the children they have educated, have expected TOC to be a relatively short-term project that will uncover what has to be understood with respect to the nature of efficient algorithms/computation, and will translate this understanding to engineering methodologies.¹ However, we still don’t have a good understanding of the nature of efficient computation. Moreover, all that we have achieved is very far from providing answers to “minimalistic” questions such as P vs. NP as well as many concrete and well defined problems that were elucidated from the attempt to understand efficient computation. This gap between the expectations and the current state of affairs has put much of the TOC community into an uncomfortable feeling.

Before evaluating the situation, we note that having ultimate goals which seem out of reach is customary in the history (and present state) of the sciences. Perhaps Mathematics and Physics provide the clearest, most outstanding examples, on a much more impressive time scale, but so do

¹The Introduction of Aho *et. al.* Report [1] provides a good illustration to our claims; see Chapter 4.

Biology, Medicine, Psychology and many others.

In judging such a state of affairs, and deciding if to quit or continue, indeed it is not sufficient that the ultimate goals are of fundamental importance. Simply waiting for the corresponding ultimate solutions means stagnation of a discipline, be it important or not. One also has to take a serious look at the research the discipline has generated in its “futile” attempts to solve these concrete problems and advance its understanding towards its ultimate goals. The question is what is the value of the insights obtained in these attempts. Again, we can only repeat our impression, substantiated in the previous section, that in its attempts to understand the nature of computation, TOC has obtained extremely important insights and thus has been the complete opposite of stagnant. Few scientific disciplines were fortunate enough to be so rich as to provide for two or three decades, at an almost yearly rate, exciting novel ideas, notions, results and further directions and problems that were not anticipated at all. This flow has not dissipated yet, and the creative forces behind it should be allowed to continue in high gear.

Another aspect of research in TOC is the depth of mathematical techniques and tools, both borrowed and invented, that were used to make more recent discoveries. It should be recognized with pride rather than disdain that the problems we have not solved yet will require more mathematical education, more knowledge and understanding of past TOC research, more hard work and longer time. Both TOC and its individual members need more stamina in research. Such stamina has always paid off and many important long standing problems (some open for 20-30 years) were solved. A few examples that jump to mind are

- The polynomial-time algorithms for linear programming and polynomial factorization.
- The linear-time algorithms for minimum spanning trees and polygon triangulations.
- The superpolynomial lower bounds on monotone circuits.
- Understanding the power of non-determinism and randomness in space-bounded machines (e.g. $NL=CoNL$ and $RL\subseteq SC$).
- The competitive k -server algorithm.

Even more importantly than solving single clearly-stated problems, stamina and the unique interaction of various subfields in TOC was responsible to the patient construction over 10 or 20 years of complete theories, starting from scratch and ending with a great deal of understanding, such as cryptography, pseudorandomness, hardness of optimization problems, and interactive computation. Algorithm design in many models (e.g. sequential, parallel, distributed and on-line) and for many types of problems (e.g., combinatorial, geometric and symbolic) gathered impressive array of techniques and paradigms, very often borrowing from advances in other subfields of TOC. Many other areas like computational learning theory, program checking and “proof complexity” are progressing fast but still have a long way to go.

It is the last area (i.e., proof complexity) that brings us back to the P vs. NP question and the other unsolved problems, which indeed may take decades to understand. It illustrates a final (perhaps unique) feature of TOC community when faced with problems on which all attacks have failed so far. This area attempts to understand formally the sets of methods that were used in lower bound proofs, and *prove* that they were not sufficient for this formidable task. In its very short existence this area already provided concrete results of this form, generated fundamental questions and novel connection of TOC to logic, and even new (unintended) insights to the practical area of automated theorem proving.

If you still feel frustrations, they must be of the healthy and natural kind that accompanies all researchers in their efforts, and drives them to continue with more energy.

3.2 Lack of a leadership group

Historically, leadership in science takes two forms. One is the traditional scientific leadership by which the prominent scientists, mainly via their own research, influence the agenda of the discipline. TOC never lacked this kind of leadership, which is clearly evidenced by the success of the field. But here we refer to the second type of leadership, the one which radiates its conviction of the importance of the discipline both inwards (especially, to the new generations) and to the outside world. We claim that TOC lacks such a leadership and that this is the source of much of its problems.

Let us confine ourselves to consider the representation and advocacy of the discipline towards the outside world. In all other disciplines this is done by the most prominent scientists in the discipline. Physics, for example, is well known to (always) have Nobel Laureates or equivalent calibre scientists in most key NSF positions. It seems that the prominent scientists of TOC never made it their business to represent and advocate TOC towards the outside. An interesting question is what is the reason for this lack of leadership. One simple answer is that it was not severely needed; as long as support for the field was arriving and jobs existed plentifully, it was natural for our leaders to concentrate on research rather than on politics. Due to our youth (as a field) and lack of experience and foresight we did not flood key positions, say in the NSF, with our Turing award winners or equivalent calibre scientists. But clearly, to effectively represent a scientific discipline, one should enjoy the respect of the community one represents, as well as have a global view and understanding of the discipline, conviction of its importance, and the ability to convey both to the outside decision makers. We were caught unprepared when money to funding agencies was drying, and the fast growth of computer science departments stopped.

Realizing this major error and acting on it can remedy the situation! Skillfully communicating the fundamental nature of the following billions-of-dollars projects, Physics lobbyists could secure government funding for the High Energy Colliders (merely one aspect of Physics research) and Biology lobbyists did likewise for the Human Genome project (merely one aspect of Biology research). There is absolutely no reason that TOC, a theoretical area that could survive a millennium on these sums of money, whose agenda is as fundamental and whose record is impeccable, cannot secure the modest sums it needs for the continuation of the excellent research within it.

Another political battleground is Computer Science itself. Both funds and faculty positions for TOC and Applied Computer Science (ACS)² come from a general Computer Science pool, and a particular effort should be directed at a just division of these resources between them. This division should be based on the fundamental contributions and achievement records of TOC and ACS to Computer Science. It requires not only a full understanding of what TOC has achieved and plan, but a similar understanding regarding the work done and planned by academic ACS. It will be revealing to read an essay like ours (i.e. one which states its foundations and gives evidence) written by leaders of ACS on the achievements and goals ACS. Even more illuminating will be an independent study (say by NSF, and perhaps also by the deans of major engineering schools) of the relative return of their investments in TOC and ACS over the last 20 years, say, judged with respect to the long term effect of the proposals and people they supported.

²For lack of better term, we refer to the part of CS academia which is not in TOC as to Applied CS (ACS). This include all faculty members in systems and applications research.

Finally, we completely agree with Aho *et. al.* Report [1] that a popularization of the main achievements of TOC is needed. However, we put this responsibility too on the senior leaders of TOC rather than on the whole community.

3.3 The outside pressures

The internal problems of TOC described in the previous subsections prevent it from effectively dealing with the outside climate which is indeed problematic. The Aho *et. al.* Report [1] mentions the saturation in the academic job market and the shrinking funds for U.S. funding agencies. We add to this list the pressure put by non-theory computer scientists and by university administrations on TOC to contribute to applied computer science projects. Fortunately for the rest of TOC, this last problem is relevant mainly in the U.S.

There is little to say about the saturation in the academic job market and the shrinking funds for U.S. funding agencies, except that these are general conditions which effect all sciences and in particular all of computer science. (However, the effect of the saturation in the academic job market is more dramatic to computer science which was accustomed to an exponential growth rate in academic positions in the previous decades.) Thus, TOC has to deal with these factors as any other discipline. To the best of our knowledge no other discipline has chosen to change focus away from its main scientific goals in view of these factors. Instead, other disciplines chose to focus on their main scientific goals and to fight for resources required for the fulfillment of these goals. The previous subsection demonstrated what we can learn from other disciplines in this regard; namely, send our best people, equipped with a strong conviction in the importance of our discipline, to this fight.

We now turn to the pressure put on TOC by Applied CS colleagues, department chairs and deans to contribute to applied computer science projects. This is indeed a unique situation that results from our membership in computer science departments (which are often in engineering schools), whose existence is motivated in part by its influence on technology. Thus, it may not suffice for us here to show that we are doing fundamental work. We also need to show that our work and teaching is relevant to the development of computer science practice and to the education of computer science engineers. We have no doubt that the case can be made that our discoveries and understandings so far had such influence, unprecedented in volume as well as in the short time it took to transform some theoretical ideas to working systems. Another important case to be made is that the education we give in theoretical courses on general modeling, analysis and problem solving abilities generally had (and will) have more lasting influence on the graduating students that went (and will go) to industry than applied courses. We must be open to the requests of ACS faculty and students for collaboration and help, but maintain the symmetry of the relationship – they cannot demand us to work for them just as we never dreamt of forcing them to work for us. Finally, chairs and deans should be convinced to judge work not by the grants it attracts (as this is already biased against TOC and by the recommendations of Aho *et. al.* Report [1] will be further biased). Rather, being academic leaders they should judge the work by academic standards, and if needed, influence funding to correct biases.

None of these tasks is easy by any means. However, we repeat that conviction in our main theses, and the willingness of our leaders to invest time and effort in these tasks is essential.

Chapter 4

Critical reading of the Aho et. al. Report

In this section we take the task of pointing out our major disagreements with the Aho *et. al.* Report [1], as well as the sources of these disagreements. This task is made hard as the report fails to explicitly state the methodological foundations on which it bases its views. This lack in itself is a major flaw. It is the source of many undefined notions of key importance to the understanding of the issues involved. It is the source of many unsubstantiated claims, even those we agree with. Indeed, it enables the report to contain contradictory views and suggestions. These allow the good hearted reader to find consolation in a particular point even if it is in sharp contrast to the main recommendations of the report. With this in mind, we list our major points of critique.

1. The report's appearance to provide a representative and authoritative view of TOC.
2. The reports biased perspective of the relation of TOC to CS and other disciplines.
3. The report's (strange) perspective on scientific prosperity.
4. The report's failure to communicate the fundamental nature of TOC.
5. The report's failure in evaluating the achievements of TOC.
6. The report's failure in evaluating the technological impact of TOC.
7. The report's unjust critique of the value structure of TOC.
8. The report's controversial view of the history of TOC.
9. What will TOC look like if the report is adopted.

We remind the reader that Aho *et. al.* Report [1] has taken upon itself to “assesses the current goals and directions of the Theory of Computing (TOC) community and suggests actions and initiatives to enhance the community's impact and productivity”.

4.1 Authority and representation

The first and foremost danger is that the Aho *et. al.* Report [1] report is taken as *the* representative and authoritative view on its subject matter (which amplifies the effect of every flaw in it). This is

certainly (by definition) the perception of those who solicited the report. It is also the perception of the authors themselves (who present it as the outcome of a year-long community-wide effort), and unfortunately of many members of the community (who have tremendous respect to some of the authors in view of their scientific leadership). The first and foremost flaw of the report is that it is neither representative nor authoritative. Strong opposing views and assessments, as well as much as the critique below, were voiced upon an oral presentation of a draft of the report to the community months before its publication (taking place in the Business Meeting of *FOCS95*). The integrity and responsibility of the authors, especially due to the voiced beliefs that the report endangers the existence of TOC, should have dictated them to take one of the following two actions. One was to integrate and fairly represent these different views, assessments and recommendations into their report, either by inviting more people to their committee or by soliciting external help. Another was to clearly claim that their report has such strong opposition and thus represents only one view within the community. We do not understand their decision to avoid both actions and note that this by itself may damage the social fabric of TOC. Still, the damage of the appearance of the report as representative can be partly repaired by the authors publishing such a disclaimer in retrospect, and by the opposition expressing its alternative views both orally and in writing.

4.2 The relation of TOC to CS and other disciplines

We feel that the Aho *et. al.* Report [1] misrepresents the relation of TOC to the rest of Computer Science (CS) as well as to other disciplines. It seems as if the Aho *et. al.* Report regards TOC as a subcontractor of the rest of CS and the other disciplines. Our view of the relationship is different. As any independent discipline, TOC should and is constantly enriched by its sensitivity to inputs from the outside. This has often led TOC to define and successfully investigate computational models and problems originating from the outside. These include models and problems motivated by technological advances in CS (parallel and distributed computing are but two examples) as well as by other disciplines (e.g., Computational Biology, Biological Computing (i.e., DNA as parallel computers), Neural Computation, Quantum Computation and Simulated Annealing). We wish to stress that we fundamentally differ from the Aho *et. al.* Report in believing that it is the internal agenda of TOC which should determine TOC's reaction to inputs from the rest of CS and other disciplines, rather than that TOC should be harnessed to satisfy the needs of the rest of CS and the other disciplines. This difference in opinion translates to several differences.

The Aho *et. al.* Report views “application-specific theory” as part of TOC. Our view is that only a minor part of “application-specific theory” is part of TOC, specifically the modeling and abstraction of new computational models and their investigation (mentioned above). The major part is the concrete fine-tuning and implementation of algorithms, and we firmly reject this as part of TOC.¹ In our opinion, “application-specific theory” falls in the domain of Applied CS (ACS²). The ACS community, using its understanding of the fundamental theory, the available technology and the desired applications, is to invent new design paradigms and construct working prototypes (that may serve as basis for a commercial product further developed and manufactured by the industry). We challenge those who disagree with us to define the border between “application-specific theory” and Applied CS.

¹Our disagreement with Aho *et. al.* Report is not merely a semantical one: augmenting the goals of TOC to encompass “application-specific theory” requires also augmenting its creative forces as otherwise this augmentation will come at the expense of the existing goals. An abstract objection is that nothing is gained by merging two disciplines with distinct characteristics into one.

²which includes systems and applications research in academia

Likewise, the application of TOC insights to the other sciences is within the responsibility of the other sciences. Other sciences, which unlike TOC do have a *sound tradition* and some self-esteem, will never allow somebody from the outside to dictate to them research directions (and in particular when to apply TOC insights). We believe, however, that TOC concepts and results will eventually influence all sciences and that this will happen by developments intrinsic to these sciences. These sciences will then seek insight to computation and if TOC will not disappear in the meanwhile it will be there to provide it.

The main obligation of TOC towards the rest of the world, is to transfer its knowledge in courses, articles and books in a form that should be available to those who want it. TOC has been doing this too and the reason it was not done even better is due to the lack of human resources. Certainly, TOC is too tiny to take on extra tasks. Of course, a free association of TOC scientists with Applied CS or other disciplines based on mutual scientific interests is a welcomed phenomena which often enriches both parties. But a forced redirection of any of these parties, contrary to their intrinsic goals, has no place in the academia.

4.3 Future Prosperity

The main recommendation of Aho *et. al.* Report [1] is the redirection of TOC scientists towards goals which are not intrinsic to TOC. To quote [1]:

In order for TOC to prosper in the coming years, it is essential to strengthen our communication with the rest of computer science and with other disciplines, and to increase our impact on key application areas.

In contrary, our main recommendation (based on the clear *scientific* foundations laid in Chapter 2) is for the TOC community to concentrate on TOC research. Searching [1] for a justification of its main recommendation, we only found the following *opportunistic* justification (in the same paragraph):

If we do so then funding will flow to TOC from a variety of sources, young theoretical computer scientists will have good employment opportunities, our field will contribute to the revolutionary developments that are surely coming in the field of computing, and will draw the intellectual stimulation from challenging problems that demand solution. If instead we turn inward then our chances for adequate funding and employment opportunities will be compromised, we will fail to take advantage of the intellectual stimulation that comes from exposure to concrete problems, and the applications that need our help will not get it.

This remarkable paragraph, that encompasses most of the spirit and danger of the entire report, is a blend of admitted opportunism, naivity, methodological faults and lack of appreciation (to the point of contempt) for the importance and achievements of TOC. It calls for a detailed analysis.

- **prosperity, funding, jobs.** The report calls TOC to recognize that material gains lies in a non-intrinsic direction, and further to seize this opportunity. There is no *serious* attempt to assess the intrinsic goals of the discipline and to derive recommendations out of this basis.
- **turn inward.** What a compact way to say “concentrate on our fundamental research, while remaining sensitive to inputs from the outside world, and in doing so continue the TOC’s great tradition of revolutionizing the way the world thinks about computing.”

- **If we do so..., If we instead...** The authors forget that in this report, invited by funding agencies and read by deans and department chairs, they *influence* the priorities for funding and jobs for TOC, not only *predict* them. Furthermore, we fear that the *negative* influence on funding and jobs (which seem to be the report’s main concern) will be the only effect of the Aho *et. al.* Report [1]. We believe that a serious effort should be made to reverse the self-fulfilling prophecies of the Aho *et. al.* Report.
- **draw intellectual stimulation.** This must be a joke. The last thing TOC lacks is intellectual stimulation. Being in a joking mood, how about Socrates sent to learn Medicine? Beethoven sent to a military academy? Einstein sent to learn Mechanical Engineering? Karp sent to learn how to take the blood pressure of frogs?
- **demand solutions, need our help.** We oppose using newspapers headline style. Converting all TOC scientists into (hopefully good) engineers will not even solve one percent of these engineering problems. not to mention that this will dry out the source of new technological revolutions. Instead, we would suggest that those who need to solve challenging engineering problems obtain a better education in TOC. We are at their service, in most engineering schools, willing to provide in our courses some of the knowledge we have achieved and will achieve given the freedom to concentrate on our research. But we are not their subcontractors nor should we be (and nor will they want us to be if they were explained the above).

4.4 Communicating the fundamental nature of TOC

The Aho *et. al.* Report [1] fails to communicate (despite its promise to do so in its title and abstract) the scientific goals of TOC. Further, it neglects to clearly state and explain the fundamental nature and importance of the formative questions of TOC; that is, the questions relating to the nature of efficient computation. We see no justification or explanation for this failure. Furthermore, the Aho *et. al.* Report [1] fails to communicate that the TOC agenda constitutes a scientific task of gigantic proportions which is in need of the best creative forces and external support to continue its successful work. It is ironic that the Aho *et. al.* Report [1], in pointing to the need for TOC by computer and other sciences, misses this as key indication to TOC fundamental importance and impact.

In general, the scientific perspective on concepts and achievements of TOC is rarely and briefly presented in Aho *et. al.* Report [1], and the focus is mostly on the technological and industrial impacts. Clearly such a basis cannot suffice to convince anyone that a scientific discipline is engaged in fundamental work, nor to evaluate its real achievements. It plays directly to the hands of those who represent the thirst of the populous for fast technological gains, and provides them with complete legitimacy to cut support to our basic research.

4.5 Evaluating the achievements of TOC

Recall that the task the report took upon itself was to assess TOC for both the outside world and its own community, in a unique authoritative and representative capacity. It had the opportunity and responsibility, regardless of its views on the future, to do just credit to our achievements. Instead, in section 2, the report telegraphically lists (important) contributions of TOC in random order, repeating some twice (one item reads “basic data structures” and the next “advanced data structures”, and then we have “public key cryptography” and “RSA and other cryptosystems”)

in a way that would surely look to an outside reader if as the authors were having a hard time coming up with examples. In none of these is any attempt made to explain what was the conceptual breakthrough behind the term, what consequences did it have on TOC goals, science and human thoughts. Furthermore, the report fails to stress that many of these conceptual breakthrough have a fundamental message which goes much beyond computer systems. Even the detailed examples (except NP completeness) are praised mainly from their application viewpoint and even this is done from a strange perspective (e.g., “..ten billion dollar industry..”)

The authors of [1] are surely aware of the generally accepted foundations of assessing scientific disciplines that we recalled in Sections 2.1 and 2.2 above. Nevertheless, even when they praise TOC achievements, there is never a scientific argument to substantiate it. Would anyone take their word that these were indeed important? Was it so hard to do in a year? As we have demonstrated in Section 2.3, which took us 2 hours flat to compile, our field was blessed with discoveries of not only importance and depth, but whose essence and far-reaching effects can be at least sketched in a few lines. These clearly convey the fundamental nature of our scientific goals. This neglect in Aho *et. al.* Report [1] will cause significant damage.

The following example may demonstrate that what is so wrong in the report. The Aho *et. al.* Report **explicitly states** that the discovery of the Interior Point Algorithm by Karmarkar in 1984 is *equally important* to the implementation of the algorithm by Karmarkar in the following years [1, P. 6]. We hope that the authors of [1] do not really think so and that they were merely carried away in their attempt to encourage the implementation of novel algorithms. Still, the message sent by this way of evaluating scientific progress (and the strange ad-hoc justification given that follows it) is wrong and dangerous.

Beyond the assessment of the achievements of TOC, we looked for the assessment of its success. We recall that by success we refer to the rate of progress done in the discipline towards its goals. The little that is said with respect to the dating of achievements is negative. One gets the repeated impression that the most important contributions occurred in the distant past. There is no mention³, let alone stress, of the fundamental discoveries that are but few years old, their importance to our understanding, and their effects on future research. There is no mention, let alone stress, of the time, stamina, interaction and connections between distinct subareas, which made some of the most celebrated recent (and past) results possible. Lacking these, the Aho *et. al.* Report [1] creates an impression that the field is stagnant. This impression is intensified by the statement that our field has discovered hard problems which may take decades to solve. Many of these were discovered long ago, and are part of our long term goals, but while unsolved, have driven us to many of our successes.

4.6 Evaluating the technological impact of TOC

Again, we reject the current technological impact as a measure for evaluating a scientific discipline. However, if Aho *et. al.* Report [1] chose this measure, it should have done a better job using it, not to mention trying to justify its use. We claim that the Aho *et. al.* Report [1] greatly underestimate of the technological impact of TOC. All the report contains in this regard are very few (good) examples in the introduction. The first major thing which is missing is an evaluation of the grand impact of the modeling, conceptualization, unification, and problem-solving and analysis abilities, developed in TOC, on essentially every computing product. This is too often taken for granted, especially by those who like to bash theory. They should be constantly reminded of the conceptual

³Except in the telegraphic list of [1, Sec. 2] mentioned above.

frameworks and algorithmic insights they have borrowed and continue to borrow from TOC. The Aho *et. al.* Report should have highlighted this reminder. The Aho *et. al.* Report [1] should have elaborated on why are the theory courses (teaching these frameworks and insights) required from every undergraduate in Computer Science departments. Indeed why does Applied CS need us so badly, if not for these frameworks and insights. Unfortunately, Applied CS tend not to realize that also it needs TOC to further recreate and develops these frameworks and insights. The Aho *et. al.* Report [1] should have done this explanation too.

Another major point not communicated by Aho *et. al.* Report is that more often than not the technological impact of TOC ideas materialized in very roundabout ways, and stemmed from models that had little if anything to do with practical applications. There are many excellent examples available, not to mention that this phenomena coincides with common sense. Furthermore, the Aho *et. al.* Report should have pointed out that the unique interactive atmosphere in TOC often enabled amazing connections, especially between the “algorithmic” and “complexity” parts of TOC. Many of these, and cryptography is but one example, had an impressive impact on technology, which is still growing at fantastic speed.

Indeed, if seeking to seriously judge the impact of TOC on technology the authors of Aho *et. al.* Report [1] should have studied computer technology in depth, in a fair attempt to understand the academic sources of major successes. They should have studied the effects of all parts of computer science academia, as well as the ability of industry to use and integrate the outputs and ideas of these into technological progress. Without such serious study, some claims of the uselessness of many “theoretical” results is totally unsubstantiated, Similarly, their call to those who generate them to engage in fine tuning of such results to current technological needs has no justified basis.

4.7 The value structure of TOC

One of the repeating bashing targets of Aho *et. al.* Report [1] is the “wrong” value structure of TOC. In their definition of Application Specific Theory, the Aho *et. al.* Report [1] says that “.. this area has not flourished to the same extent as Core Theory and Fundamental Algorithms. This is partly due to the value structure prevalent in the TOC community which places excessive weight on mathematical depth and elegance, and attaches too little importance to the genuine links between theory and concrete problems.” The report repeats this in other places, such as after giving their examples on bridges between theory and applications, saying “In each of these cases, the exploitation of theoretical results was far from trivial, and required a deep understanding of both the underlying theory and application domain. Such applied studies require theoreticians to cope with the messiness of the real world, often at the expense of elegance, mathematical beauty and universality that the TOC prizes so much”. These strong statements are made with no qualifications, giving them the aura of universal truth.

The harm of such statements, which will be taken as guidelines by funding agencies, is enormous. The source again is methodological confusion and the neglect to distinguish between the intrinsic goals of TOC, and the goals of applied research. While there is no doubt that implementation has to take into account many factors when trying to create a (potentially) working product, the progress of theoretical work *depends* on the ability to abstract away such factors and elucidating the important ones. If Aho *et. al.* Report [1] wanted to voice the opinion that this “wrong” value structure was the reason that TOC did not *engage* in implementation and fine-tuning, they should have stressed that it is nevertheless the “absolutely right” value structure for achieving the intrinsic goals of TOC. They should have clarified that adding implementations to the responsibilities of TOC (which they recommend and we reject) generates two distinct sets of goals, which should

be pursued and evaluated according to different sets of values! It is clear that this confusion will cause their message to be used against the best proposals of TOC research, including the ones that Aho *et. al.* Report [1] concedes should go on.

4.8 History according to the Aho *et. al.* Report

The Formative Years. The Aho *et. al.* Report identifies 1955–75 as the “formative years” of TOC. It is not clear what is meant by this term and what are the criteria for identifying the *formative period* of a discipline. We are not even convinced that this term can be given a meaningful definition. What is the formative period of Physics? Which years would have physicists of *various times in history* identify as the formative period of Physics? We of course completely agree that the years 1955-75 were crucial in determining the major directions of TOC, but the same holds for the next 20 years.

Maturity. The Aho *et. al.* Report state as a fact that TOC has achieved maturity (e.g., see [1, P. 11]). Again, it is not clear what criteria has been applied here and our personal feeling is that TOC is still in its childhood. When progress in the field is currently proceeding in full speed, where is the sense of maturity emanating from? This point of view serves the recommendations of Aho *et. al.* Report [1], as it creates the atmosphere that there are less things to do in TOC, and we can turn to other objectives such as implementations and fine-tuning. Again, no justification is given by the authors to their evaluation of maturity, and we completely reject its implied consequences.

We are not professional historians, and our laymen feeling is that this kind of historical judgments are better left to times in which one has a better perspective of the past. Certainly, this should be true in an active discipline which constantly generates new notions, connections and targets for study in its strive to understand efficient computation. If the Aho *et. al.* Report [1] still wanted to take on such a role, its authors should have justified their assessments by a serious comparative study of the evolution of TOC and other disciplines.

4.9 TOC according to the Aho *et. al.* Report

Finally, we speculate on how will TOC look like if the recommendation of Aho *et. al.* Report are adopted (especially by parties alien to TOC).

Who will join TOC? The Aho *et. al.* Report states that in the last 20 years TOC has “attracted to its ranks some of the brightest young scientists” [1, P. 2, L. 4–5]. Why were these bright scientists attracted to TOC? Will they be attracted to TOC when it takes the form advocated by the Aho *et. al.* Report [1]? We believe that these bright scientists were attracted to TOC because of the strong appeal of the fundamental questions which TOC is all about. They were attracted too by the exciting atmosphere of research and achievement TOC. While we may have less jobs to offer to the future generations, we *have* to make sure that we continue to attract the most creative and talented. The precondition of fundamental and important questions is threatened to be taken away by adopting the recommendations of Aho *et. al.* Report [1], and the precondition of exciting atmosphere is reversed by moods as portrayed in Aho *et. al.* Report [1]. Realizing the disaster of the field from reduced quality of its researchers, we have to fight to ensure that they enjoy the same conditions we had.

What will be funded and what effect will this have on TOC? When reading this, recall that funding means much more than money. It is used by departments as a measure of success in research, and thus as criterion in promotion and tenure. This is reasonable *only when* funding decisions relate to merit of proposals to the intrinsic goals of the discipline.⁴

In a private discussion, Karp has conceded that a proposal of the nature of the first work on 2-prover proof systems may not have been funded under the recommendations of the Aho *et. al.* Report [1]. It should be emphasized that such decisions must take the responsibility of giving up the unforeseeable fruits of this model, including PCPs and the revolution in understanding approximation mentioned above.

But 2-party proof systems is merely one example. We claim that the attitude reflected in the Aho *et. al.* Report discourages the introduction and development of revolutionary ideas. Such ideas tend to sound non-applicable and (especially junior) faculty will be and actually are already (directly and indirectly) discouraged from pursuing them. These revolutionary ideas are the engine of any science and were the key to so much of TOC success. Thus, the attitudes reflected in the Aho *et. al.* Report constitute a big danger to science.

The good hearted may point out that the Aho *et. al.* Report [1] does suggest to continue supporting research in “core theory and fundamental algorithms”, and not to shrink funds for this support. This is made in such minor a voice, is almost unsubstantiated, and sharply contradicts the main message regarding values (above) and important directions, and so the net effect on the decision makers is unmistakably bad. This is not a mere feeling of ours – unfortunately we already heard stories to substantiate it emanating from both NSF and computer science departments. It is urgent to fight back to undo the damage on this issue.

Will there be interaction between the areas of TOC? The Aho *et. al.* Report seems to recognize the importance of interaction between the various areas of TOC. We fear, however, that their recommendations will tear TOC apart into non-interacting entities confined in their application areas. This will be the direct consequence of Aho *et. al.* Report [1] demand from TOC researchers to invest time in implementation and fine-tuning the results in their area, which will drain the time for keeping up and interacting with other areas of TOC. We leave it to the reader to imagine the status of TOC research, if our field was so fragmented 10 or 20 years ago.

Can scientific revolutions spring out of application-specific theory? Surely they can, but at much greater effort and lower probability than from a free scientific process; that is a scientific process which evolves according to its own intrinsic logic and is not harnessed to the immediate technological needs. The Aho *et. al.* Report does acknowledge this point and seems to suggest to maintain a small avantgarde for these purposes [1, P. 3, L. 27–35]. Our point is that the TOC community is below the critical mass and cannot afford having its most brilliant minds taken away. We consider any form of discouraging TOC research to be a dangerous impediment to science.

The applicability of TOC. Finally, we get to the very thing that Aho *et. al.* Report aims at: the applicability of TOC to real life. They mention some of the past applications of TOC to computing practice, but neglect to say that very often these applications emerged out of a free scientific process which was not aimed at achieving these technological goals. Furthermore, if it was not for the work done at the foundations of TOC there would have been very little that the Aho *et. al.* Report could suggest to implement. Yes, we guess that one may say that we have

⁴Unfortunately this is not always the case.

enough foundations by now (and it is time to do real building). This reminds us of the comment “too many notes” attributed to Emperor Joseph II hearing *Cosi fan tutti*. Again, we disagree with this thesis completely.

4.10 Summary

The Aho *et. al.* Report [1] suggests that TOC harness itself to help technology and other disciplines. Our answer is that TOC is too small a field (in number of scientists) to handle the fundamental questions which constitute it. The intellectual challenges of TOC are gigantic and of the greatest importance. It is thus essential that this discipline is given a high priority, both inside CS and among the sciences. The very least that should be asked is not to disturb its successful momentum.

Chapter 5

Our Recommendations

The basis of all our recommendations is our main theses (abbreviated below), which we hope to have justified in our essay.

Thesis: *TOC is a fundamental scientific discipline, which has achieved tremendous productivity and impact. This success of TOC continues to this very day, and is likely to continue and grow. The source of this success is the ability of TOC to attract the best creative minds into it, and the ability of these researchers to pursue the intrinsic goals of their field with full academic freedom.*

Our recommendations are:

1. **Recommendation to TOC – do good research in TOC:** Theoretical Computer Scientists should concentrate their research efforts in the Theory of Computing while continuing the tradition of sensitivity to *scientifically relevant* outside inputs (e.g., from the rest of Computer Science and other sciences). Each TOC scientist should determine his/her research directions in TOC based on his/her own understanding of the intrinsic goals and current state of the TOC.
2. **Recommendation to TOC – strengthen self-esteem:** TOC scientists should realize the fundamental importance of their discipline and its achievements, further their conviction to their colleagues and defend TOC from outside pressures.
3. **Recommendation to outside bodies (universities):** Academic institutions (mainly deans and department chairs) should study and verify our thesis. They should continue the two-century-old tradition of allowing faculty to pursue their intrinsic research interests while enjoying full academic freedom. In allocating positions between (sub)disciplines the relative importance and success of the disciplines (as defined in Chapter 2) should serve as the guideline. In general, decisions should be based on expert evaluation of the contribution of candidates to the intrinsic goals of their field.
4. **Recommendation to outside bodies (funding agencies):** Funding agencies should study and verify our thesis, and continue to supply the funds required for the continuation of TOC research at the current momentum. In dividing funding between different disciplines, the guideline should be the relative importance and success of the disciplines (as defined in Chapter 2) scaled by the cost of research in each discipline. Peer review, based on the potential scientific contribution to the intrinsic goals of TOC, is by far the best way to evaluate TOC proposals for funding.

5. **Recommendation to senior TOC – support good work in TOC:** Senior members of the TOC community should continue the tradition of supporting any good work in the TOC, regardless of its immediate applications, as long as it is governed by a candid desire to advance the TOC.
6. **Recommendation to senior TOC – representation in decision bodies:** Senior members of the TOC community must take it upon themselves to communicate the importance and success of TOC to the decision making bodies of academic institutions and funding agencies. This in particular requires the senior member of TOC to represent TOC in these bodies, and to participate in forums which influence their decisions.
7. **Recommendation to senior TOC – communication to the outside:** Senior members of the TOC community must take it upon themselves to communicate the importance and success of TOC to their colleagues in Computer Science departments, other scientists, and the general population. This should be done in conversations, lectures, expository articles and books, and any other form they can think of.
8. **Recommendation to senior TOC – communication to the inside:** Senior members of the TOC community must also take it upon themselves to communicate the importance, success, history and impact of TOC to the junior TOC generations. This should be done mainly via their curriculum courses and textbooks, but also in seminar lectures and private conversations.
9. **Recommendation to junior TOC – study TOC in depth:** Junior TOC scientists (and in particular graduate students) should study and deeply understand the scientific goals of their field, and the record of its successful progress towards these goals. This will result in higher self-esteem and ability to confront the many outside pressures awaiting their career, as well as help them to do even better research.
10. **Recommendation to all of TOC – study history of science:** Studying the philosophy, history and sociology of Science may further enrich the perspective on the goals of TOC and its current status. It may also teach us the ways in which scientific disciplines coped with various problems in their history.
11. **Recommendation to Applied Computer Science – learn TOC:** Applied Computer Science students should take courses in TOC. This should be done at a level allowing them to make good use of TOC research. In particular, they should be able to take a reasonably written TOC paper and implement an algorithm presented in it as well as do the necessary fine tuning. This will allow them to cope with the task of transporting theoretical frameworks and insights to application, a task which is within their domain of responsibility.
12. **Recommendation to Applied Computer Science – teach TOC:** Applied computer scientists in the academia should make available to TOC and other interested parties their own perspective of the goals, achievements and directions of their discipline.

Finally, we call upon readers of our essay who disagrees with our thesis or any other part of it, to articulate their disagreement. An open and free debate of these issues is the best way to understand how to guarantee that the future of TOC is as successful as its past.

Bibliography

- [1] A.V. Aho, D.S. Johnson, R.M. Karp (Chair), S.R. Kosaraju, C.C. McGeoch, C.H. Papadimitriou, P. Pevzner, “Theory of Computing: Goals and Directions”, University of Washington Technical Report CSE-96-03-03, March 15, 1996. Available at the following URL: <ftp://ftp.cs.washington.edu/tr/1996/03/UW-CSE-96-03-03.PS.Z>.
- [2] C.H. Papadimitriou, lecture in the workshop in honor of Karp’s 60th Birthday, *FOCS95*.
- [3] W. Shakepeare, “Julius Caesar” (Act 3).