Research Evaluation Exercise

Dipartimento di Informatica
University of Pisa

August 1999

This report is the result of collaborative work by a panel of computer scientists with a variety of scientific backgrounds, active in diverse research institutions as well as different parts of the world. The panel included:

Martìn Abadi, Bell Labs Research, Lucent, Palo Alto, CA, USA
Gilles Kahn, INRIA, Rocquencourt, France
David Notkin, University of Washington, Seattle, WA, USA
Franco Preparata, Brown University, Providence, RI, USA
Jeffrey Ullman, Stanford University, Stanford, CA, USA

The work of the panel of reviewers is based on three components:

1. *Written Documents*. Before the reviewers met, they received the Annual Research Report of the Computer Science Department, published in January 1999. This report on research and teaching contains an exhaustive presentation of the research carried out in the department. It is organized by themes, with a short presentation of the background and the research objectives of each theme, as well as a detailed list of personnel and a complete list of publications. The reviewers have also received a document containing 14 research proposals, organized along a common structure: State of the Art and Trends, Relevant Research Activities at the Department, Short Term Plans and Expected Results, Long Term Scenarios. Each proposal also included short CV's of the participants, essential references of the field, and main publications by the research team.
2. *Oral presentations, discussions*. The reviewers attended a meeting, convened in Centro S. Paolo all'Orto, Pisa, July 8-9, 1999. During this public meeting, after a brief and useful introduction by Franco Turini, Department Chairman, all proposals were presented, and there was time for questions and some private discussions between the reviewers and the research groups. Rector Luciano Modica provided additional context for the role of research evaluation. A brief presentation of the reviewers' first impressions was given at the end of the meeting.
3. *Electronic mail exchanges*. During July and August 1999, the reviewers exchanged their notes electronically, then drafted and discussed sections of this report. Gilles Kahn served as a coordinator during this phase. Additional information on the research work of the Department was also found on the World Wide Web.

It is apparent that the research evaluation exercise is taken very seriously by the Department and that its planning has involved a large effort. The research proposals are far from reproducing the Annual Report and put forward substantial modifications. The reviewers are conscious that this event, having scientists from outside the University and even outside Italy evaluate the Department's research, is novel and unique. They are honored to be part of this process and they hope that their observations will be useful.

The report is divided into two parts. The first part discusses general issues for the Department, such as the overall balance of the research topics, the composition of the faculty, or the connections to other disciplines. The second part of the report contains an individual discussion of each one of the 14 research projects that were presented, with comments on strengths and weaknesses, and often constructive suggestions for the future. The projects are evaluated in comparison to related, first class international research efforts and not compared to one another.

We have felt free to include all sorts of observations. Of course, we cannot be expected to fully understand the Italian university system –which is undergoing changes anyway– in the short time devoted to the evaluation. We consider that our role is merely to provide an external view of the department's research activity and to indicate general directions for change when it is seems appropriate. As an additional *caveat*, we cannot be experts in every area covered in the department, and we have not read all the papers cited.

Part I
General Issues

1. **Context**
The Computer Science Department of the University of Pisa is well-known world wide. Several of its intellectual leaders have a prominent position in their field. The organization of this review, conducted in the most professional manner, is a confirmation of this reputation. The reviewers are impressed by the overall quality of the department, the intelligence of the presentations, as well as the open style of the exchanges with the researchers. It is clear that the department has a distinguished role in Italy, by the number and quality of its students and the level of its international scientific connections. The global publication record is excellent, with presence in journals and conferences with the highest scientific standards. Furthermore, a number of former students of the department pursue very distinguished academic careers in the world. Any criticism of the department must be understood against this background of overall excellence.

2. **Projects**
The presentation of the research activities as a collection of research projects is a good idea. This gives an opportunity to stress the focus of the groups, to discuss the combination of talents, the track record of the senior researchers and their national or international collaborators, the relation between groups, to understand the overall character and balance of the department's research, and more generally to look toward the future. The research projects make a serious attempt at discussing future research topics, and this is highly appreciated.

Sometimes, a group is too small for its ambitions, or not very coherent and it will be criticized on these grounds. Still, it is better to have one's research plans discussed than not presented at all. A number of faculty, whose name are found in the Annual Report, do not seem to be associated with any of the 14 research projects. The reviewers do not know the meaning of this absence, but it is a fact that the corresponding research is not discussed, possibly leading to a slight imbalance in the reviewers' global view of the department.

3. **Research Topics**

There is no doubt that the department is heavily oriented toward theoretical work. This has many advantages: it attracts high-powered students, it doesn't need heavy experimental equipment, it gives an opportunity to have contacts with highly skilled individuals around the world, it allows the young researchers to focus on the single activity of publishing papers in conferences and journals. There are also drawbacks of course, such as rejecting creative students who are not theoretically inclined, or being less directly aware of the actual problems raised in the rapid evolution of digital technology.

In terms of Computer Science theory, the department is clearly focused on logics and semantics. This is useful in many ways, in particular for teaching fundamental concepts of programming, programming languages and distributed systems in an efficient manner. But this focus, we consider, is too systematic. For example, we feel that there are far too many people with a logic programming background, leading at times to an impression of a single culture. There are many theoretical areas that are not covered at all, and there must be consequences in teaching. To be concrete, here is a short list of theoretical topics that are apparently not covered at all, or covered only lightly, and that are relevant for the future:

Computational Geometry
Theorem Proving
Cryptography
Probabilistic Modeling of Computer Networks
Symbolic Algebra
Computational Linguistics
Information Retrieval
Image Analysis and Synthesis

This list serves to show that an increase in the diversity of topics doesn't imply a reversal of the fundamental theoretical orientation of the department. We are not suggesting that the department pursue each and every one of these topics, but rather we wish to point out that even in theoretical topics, the department could be broader.

Along with this impression of excessive emphasis on logics and semantics at the expense of other areas, there is, to outsiders like the reviewers, a very surprising level of recruitment by the department of its own graduates. We have not tried to do a detailed accounting, but this phenomenon, whatever its causes may be, is blatant to a really unusual extent. Self-recruitment and scientific monoculture go hand in hand, and the department will have to make a conscious and strenuous effort, in the next few years, to steer away from a course that would impoverish it intellectually in the long run. The recruiting system for faculty seems to have changed recently, which may provide the appropriate opportunity for action.

4. **Broadening Computer Science Research**

Another point has an impact at this stage. In the last five years, computing technology has been changing at an incredible rate, and there are no signs that this rate will slow down anytime soon. Few project presentations discuss this background explicitly. The main changes concern (1) the spreading of computer technology to an ever wider spectrum of industries (2) the commonplace "convergence" between telecommunications, computing and audiovisual technology, and the scale of distributed computing with the Internet and the World Wide Web.

As a result, there is an increasing and widening demand for students with a Computer Science education in many parts of the world. The last analysis showing the first

employment of Pisa graduates dates back five years and it would be a good idea to update it and then maintain it continuously. In terms of research also, it seems that understanding better where students are employed would give an incentive to prefer new directions that correspond to a higher long term demand. As an example, other Computer Science departments have experienced an explosive demand from the telecommunications sector, and right or wrong, it has not happened for Pisa Computer Science.

Basically, all Pisa Computer Science students who obtain a Ph.D. embark on an academic career, in Italy or elsewhere. This is not the case in several countries we are familiar with, and it may explain in part the relatively small number of Ph.D. students. This situation is not ideal in the long term for many reasons, in particular because it has an impact on research directions. When a student is undecided about his/her own professional future, a normal strategy is to keep options open and choose a thesis topic that is both intellectually challenging and somewhat appealing for an industry employer. There is a risk, if only an academic career is contemplated, that topics allowing the production of a maximum of papers be preferred. Theory has pitfalls too, and in Pisa some projects think more in terms of finding all the possible applications of the theoretical tools they master than in terms of problems that need to be solved.

We feel that there should be tighter intellectual connections with real problems. More frequent partnerships with industry would make it possible to be present when a problem is born, to participate in defining the future problems and not only their solutions. It is not easy to develop industrial partnerships, but Italy is an advanced country with many leading high-tech industries that are eager to hire bright students.

There is very convincing work on applications in this department and it is important that "experimental" or "applied" should not mean inferior, ad hoc or without a scientific merit. Of course, we are not advocating dropping good theory work in favor of mediocre experimental work, but we argue for a greater variety of research attitudes, for an atmosphere of research where all kinds of projects can be attempted. To monitor this, in the medium to long term, the department could set itself some targets for its research makeup, deciding for example on a percentage of effort for theory, core technology (e.g., networks) and applications.

Several projects cooperate with other departments of the university, with the Santa Anna school or the CNR. This is positive, but we think that these collaborations could be even more frequent, given the wealth of possible interactions of Computer Science with mathematics, biology, computer engineering, etc. It is clear that there are some institutional barriers to cooperation, but it should be an objective of the department to overcome them in the interest of research.

5. Evaluation and departmental strategy

No serious change will take place unless the evaluation of research is given some consideration. It is clear that the current evaluation system emphasizes publications. This is good, and our recommendation is in no way to discount the importance of publishing. But publishing is a way of making one's research known, and Computer Science research has a variety of possible outcomes: novel ideas, novel pieces of software, systems, patents, even start-up companies. What counts, whatever the outcome of research, is its creativity, its originality, the surprise that it contains for the fellow scientists, and also its impact on the outside world. In this regard, while the department is publishing in highly respected forums, these tend to be in select communities and to have a modest impact. Research evaluation should progress to incorporate a more global view of research, that focuses on originality in general, whatever its means of expression, and not solely on publications.

The discussion above shows that, while the department is quite successful in research, there is a need to organize some strategic thinking. In our eyes, it is certainly a merit of the review process that it helps in pointing out a number of questions that deserve discussion and decision for the future. The department is intellectually very well positioned, it has a bright and professional faculty. It is highly regarded in Italy and in the world. It has taken an innovative course in having an external evaluation. We hope that our remarks will be useful in planning its future strategy.

Part 2
Research projects

Project number 1.   Rigorous Methods for Requirements Capture and Software Architecture - Vincenzo Ambriola, Egon Börger, Carlo Montangero

The proposal concerns methods for capturing requirements on software, for developing architectural models, and for refining those models through formal methods. The requirements capture starts from informal statements and yields formal ones, relying on "shallow parsing" and user feedback. The subsequent formal work is based on Abstract State Machines (ASMs) and the logic Oikos-adtl. The focus of the project is practical: the oral presentation emphasized "theory for an application, not for its own sake".

The participants have sufficient expertise in the research areas represented in this project. Moreover, in different communities and to a varying degree, they have international visibility, many publications in conferences and journals, and several helpful contacts, for example at Siemens and Microsoft.

The association of the participants in this project is new. Although in the short term it may take significant effort on all sides to create a powerful, coherent project, this association could be fruitful.

Thus, the participants are in a good position for working toward the goals of the project. Moreover, those goals are quite worthwhile. In addition to developing general techniques and tools, the project may address interesting examples, sometimes in collaboration with industry; the treatment of those examples could be useful in its own right.

However, we are somewhat skeptical about aspects of the project. The researchers face a number of challenges in trying to achieve the stated goals.

First, it is generally quite difficult to make progress on practical software engineering problems through the application of formalisms. It can and does happen, usually when applying the formalism to narrow problems. The proposed, broad application of ASMs is quite ambitious and its benefits should be continuously evaluated during the project.

Second, it is also quite difficult to develop effective general-purpose solutions to software engineering problems. More specific efforts, focused on a particular domain or class of systems, are often more successful. For instance, it may be extremely difficult to use "shallow parsing" of natural language for subtle specifications like that of the Java VM (an example mentioned in the presentations), but "shallow parsing" may work well for other kinds of specifications.

Third, in requirements elicitation, it is crucial and difficult to find the right balance between what a human does and what the system does. The difficulty is compounded by the iterative nature of requirements elicitation. For this project, the difficulty will appear in the process of going from natural language descriptions to ASMs. (As a minor point, a distinction was made between "skilled" people – who can read ASMs – and "unskilled" people – who cannot. Collaboration with industry will only be harmed by using language in this manner; domain experts are crucial to most requirements efforts and they are very skilled even though they cannot read ASMs.)

Fourth, conventional wisdom says that requirements should focus on behavior rather than structure since the latter touches on design. Since this project takes an alternative view, closely linking behavior and structure, the project should justify and validate this view. A related issue is whether one can always find a satisfactory, useful "ground model".

Finally, we encourage two specific activities for the project. The first is to define concrete plans to build artifacts that would support the project and that would use its ideas. Building such artifacts can provide focus. Artifacts are also crucial as vehicles for experimentation and, in this case, as tools for debugging specifications. Thinking now about these artifacts would be beneficial. The second is to start thinking now about how to judge the progress of this project, how to articulate its intended results, and how and where to disseminate those results. Collaborations with industrial research and also industrial development should play an important role.

Project number 2. Type Based Models and Query Languages for Web Data - Antonio Albano, Giorgio Ghelli

This strong project focuses on the topic of semi-structured data, with a particular focus on typing issues. Led by two faculty members with excellent credentials in databases and type theory, this effort represents a productive balance of theory and practice. Our expectation is that this group will produce a number of strong research results in this project.

The world wide web provides the central motivation for studying semi-structured data: at the present and in the near future, at least, it is unrealistic to expect complete schemas (and associated type information) to be provided for web-based information that would be preferable to manage as databases.

The research challenge the group is undertaking is not a small one. In particular, they need to figure out how to take advantage of their expertise in type systems in a context in which explicit types are (almost by definition) difficult to find. Regardless of the approach this group selects, it will be crucial to use the experimental component of the effort to ensure that the theoretical work provides utility and stays on track.

The researchers discuss how they can take advantage of situations where they know something about the types that can appear in a certain context in semi-structured data. This seems possible, although difficult, since the nature of semi-structured information is that you don't want to restrict structures for all time, because types could be inferred and later violated.

One suggestion is that it may be possible to build a system that maintains types dynamically, or even to build a system that enforces certain type constraints and leaves others flexible. Determining this requires that theory be coupled with practice. Rather than

developing a type system and leaving it for others to try, we suggest coupling the development of the theory with a demonstration project, and *see* whether users can deal with and benefit from the science.

Another question we had was in what dimensions will type information actually help with web-based data. The project material claims that "types help for optimization, correctness, intelligibility of data description, integration with structured data." Determining, empirically, which of these are most important (and feasible) in practice for this domain is important. The answer to this question may depend heavily on the class of applications studied, as well.

We had some minor concerns about the conjecture about ambients and about the absence of any material on query planning (which is more difficult in such a context).

Overall, we are very positive about the project. To make real progress will require significant research efforts, both theoretical and experimental; we feel that this group is well-positioned to succeed in both theory and experimental dimensions, as well as in the connection between the two approaches.

Project number 3. Web Computing and Adaptive Agents - Giuseppe Attardi, Maria Simi

This project represents the work of two faculty, one with a background in networks and distributed computation, and the other in traditional AI. The objective of the project is in a sense a combination of the two viewpoints: develop sophisticated Web applications. The reviewers feel that, while the project is not without flaws, it represents an excellent direction for the department to take. We encourage the department to grow this effort if at all possible.

The faculty involved have a record of building some useful and clever applications. For example, a garbage collector they developed was adopted at SUN. They have built a form-filling interface for tax forms, an agent for helping students with their study plans, and have experimented with search engines technology, including the use of natural-language to support Web-page classification. In addition, Giuseppe Attardi plays a key role in national networking infrastructure.

Altogether, the project is in a fun, lively area, that should be very attractive to students. We imagine that students get exposure to both modern networking ideas such as CORBA, and to some of the popular AI technologies such as agents and natural-language.

Hence, it is surprising, given the nature of the research goals, that the principals appear to be having trouble attracting students. There seem not to be Ph.D. students working with them, and all that they have accomplished is with the involvement of undergraduates. We are unsure of the cause for the situation, and we wonder whether the problem lies with admissions policies. Are students with an applied interest being rejected for the Ph.D. program, or perhaps going into Engineering and becoming unavailable to these faculty?

We also have some concerns about the environment in which this research is being conducted. The Web is an unruly place. It may be premature to invest heavily in theorizing about it; we need more experimental knowledge of what is possible, what works and what fails. At one level, the project is an attempt to understand the Web by building sophisticated applications, and this aspect would be right at home in Silicon Valley.

However, perhaps in deference to traditional academic values, the project also speaks of developing theories in a number of areas, including distributed object models, intelligent agents, and learning. Perhaps most incongruously, they are publishing papers in journals about their experiences on the Web.

The notion of what is great software research has changed over the years, and this project is trying to follow the trend. However, we see roadblocks, some of which may be due to the traditional measures that are applied to research. We wonder whether they receive encouragement from their colleagues to build, and teach their students to build, novel applications rather than writing journal papers about how they imagine such systems might be built. We believe that if they follow the experimental road, the ideas will come, and others will hear of their achievements as long as they are conscientious about describing their ideas on-line and perhaps at conferences and workshops.

While we are highly sympathetic to their problems, we also share a feeling that there are ways the project could become more effective. Most importantly, there is the sense that it is attacking too many problems at the same time. As a result, a central theme for the project is missing, and the project gives the appearance of having too scattered a set of goals. Further, these goals cover a wide range of levels of abstraction, from low-level integration mechanisms like CORBA to highly speculative and multifaceted software architectures, such as "intelligent agents."

We acknowledge that a collection of short-term projects is appropriate for undergraduate theses. For senior researchers such as those involved in this project, and for first-rate Ph.D. theses that we hope will come out of this project, it would be necessary to focus the project more. This focus has two aspects. First, it would be useful to pick a long-term (i.e., long by Web standards, perhaps three years) goal for an exciting and innovative application. Second, we advise selecting a *small* number of technologies that the project will try to advance, in the course of creating the application.

We are heartened to see an enthusiastic attempt to do serious applications. We hope to see the project supported not only with money and equipment, but by the creation of an environment where such applied software developments are appreciated by colleagues and potential students alike. A large fraction of the future of Computer-Science research is reflected by this effort.

Project number 4. Scheduling and synchronization in distributed systems - Maurizio Bonuccelli

This is a recently constituted research group, arising from the confluence of interests of members of the department, who were previously independently operating. The objective of the project is to improve the utilization of computer-communication systems.

In this general area, which represents one of the hottest topics throughout the industrialized as well as the developing world, two specific problems are addressed: packet scheduling and clock synchronization. Whereas the former directly concerns utilization, the latter relates more to reliability and operability. Both topics obviously bear on the subject, and the investigators should strive to organically weld them in the development of their program.

On a technical level, the model chosen, after filtering inessential physical details, is meant to cover a broad variety of state-of-the-art distributed systems, and reveals adequate

awareness of the field. The communication media considered, (comprising optical fiber networks, LAN's, NOW's, mobile networks, fixed segments of phone systems, etc.) also reveal awareness of the available hardware. The standing of this group is best expressed by their publishing in most qualified venues of their field, such as the IEEE Transactions on Communications, the ACM-IEEE Transactions on Networking, and Distributed Computing. A tangible sign of recognition is that M. Bonuccelli has been Program Vice-chair of MOBICOM 97, the premier international conference on networking and mobile computing.

The research approach is basically algorithmic, and involves a variety of difficult combinatorial optimization issues. The research themes compose a relevant research plan, at least on an analytical level, and one word of caution must be said on its perhaps excessive breadth.

We feel, however, that the talent of this group will be best developed and realized by enlarging the horizon of the project to include technical milieus closer to actual implementations. This research area is perhaps not common in a department of Computer Science, especially one with a marked theoretical orientation; it is more natural in an engineering environment. But Computer Science has an overbearing engineering component, even more so as applications call for a more vigorous Computer Science participation.

For these reasons, notwithstanding allusions by the researchers to difficulties due to university organization (science vs. engineering) and to administrative chores, the panel strongly encourages the investigators (with full departmental support) to seek genuine collaborative overtures with their engineering counterpart, in order to achieve a critical mass consistent with the scope of the research plan, but, equally important, to form working partnerships with operating industrial or service concerns.

The use of real data to validate the analyses, and the verification of the soundness of the methods in real-world settings, are today nonnegotiable elements to provide full credibility to such research. Today we are witnessing a compelling convergence of computing, networking, and communications. If the department chooses to play a role in the rapidly evolving research landscape, this is clearly a strategic area and appropriate steps should be undertaken to foster its growth.

Project number 5. Security for Concurrent and Mobile Processes - Pierpaolo Degano

This project concerns the analysis of security properties. Its focus is on protocols and mobile agents. Its tools are process calculi (like the pi-calculus), their operational semantics, and static-analysis techniques for these calculi.

The project is fairly new and small. It represents the recent research program of Pierpaolo Degano and his collaborators (local and remote). Degano brings to bear his intelligence, his knowledge of process calculi and semantics, and his experience. He has a strong record of publications in international conferences and journals; the first results of this project are already leading to new publications in appropriate venues.

The focus of the project, on an important but limited set of problems, is advantageous. The project may be able to obtain interesting, useful results about particular examples or special-purpose methods, with a reasonable, bounded investment in general techniques.

Thus, the strategy of this project seems good. It might still be refined, in particular with the aim of expanding the impact of the project. A few suggestions are below.

The project draws on the results of other projects on similar topics and on a substantial body of theory, quite competently. On the other hand, more direct knowledge of security problems and systems may help in ensuring the relevance of this research. So the project may benefit from further contacts with the "real world" of security, at the university (perhaps with the engineering department?) or in industry (through research labs?).

This proposal has some overlap with that of Montanari et al., which mentions the use of HD-automata for protocol analysis. Some collaboration might be fruitful.

The project may wish to consider a broader perspective on security, for example taking into account the weaknesses of implementations. In addition, it may be worthwhile to evaluate the performance of the techniques developed with respect to accuracy and scope (e.g., do they give fewer false negatives than others? can they handle more protocols?).

Project number 6. Transportation and Logistics: Models and Algorithms - Giorgio Gallo, Stefano Pallottino

This group works on problems of Operations Research such as network flows, scheduling and non differentiable optimization. Unity in the project comes from a focus on applications to the transportation area, and to logistics in general. The group, comprising 4 professors, 1 post-doc, and 2 students has a good size and is clearly aware of the progress of the field, be it in academic research or in commercial software. The publication record is good, with more of an emphasis on journals than conferences, as is the case with this field.

Some problems that are worked on are not new. It is very pleasant to see that the solutions proposed compare favorably with competing approaches, and that there are systematic performance experiments. The evaluators appreciate the focus on concrete problems (such as bus driver scheduling) and the fact that the project has actually been working with a consumer of the technology, who is reported to be pleased with the results. There are good connections with industry, in particular with bus companies in Italy, and consultants on transportation problems. This is essential if one wants to have adequate data, of course. The researchers produce software (which is not distributed by them) on shortest paths algorithms, min-cost, multi-commodity codes etc.

The research group is clearly on the side of continuous optimization in spite of a stated interest in heuristics. It is surprising that in a Department with a heavy emphasis on logic programming, and historical competence in constraint solving, there is not more work on combinatorial optimization, which finds a wide number of applications in the transportation sector.

The group is serious and competent, reasonably coherent, but not necessarily very revolutionary. Continuous optimization is based on the ability of computing derivatives, which is often difficult, in particular when the function to optimize is expressed in a highly refined program that computes it. Hence one might expect to find people interested in automatic differentiation research in this project, or at least with experience in using experimental tools of that kind. Another query that we have is the concentration on transportation networks. If the project wanted to work more on telecommunication networks, synergy with other groups, for example Bonucelli's might emerge. In fact,

reusing techniques from transportation networks in communication networks might prove quite useful.

In terms of transportation networks, there are also a variety of new problems coming from a more intensive use of information technology in vehicles or road systems: for example, the delivery of real-time traffic data to cars on the road will allow drivers to independently decide how to change their route in response to that data.

A reviewer also suggested that the researchers should look at the recent work (primarily by Kautz and Selman) about solving hard problems using general-purpose satisfiability solvers (like WalkSAT). In many areas, these authors have shown that the conventional wisdom of taking domain-knowledge into account in a solver is **not** better than casting the problem as satisfiability and then solving it with a very fast general purpose solver.

The remarks above indicate that there are many possible bridges between Operations Research and Computer Science. Strategically, it is probably a good idea to prefer topics that will strengthen the relationship between this project and the rest of the Department.

Project number 7. Systematic Development of Semantics, Static Analysis and Verification Techniques by Abstract Interpretation - Giorgio Levi

This project plans to use abstract interpretation to a variety of problems in semantics, analysis and verification. The title and the presentation show very clearly that the promoters of the project want to distance themselves from logic programming, even though this is everyone's scientific background and basis for an international reputation. A priori, this is to be seen very favorably, in view of a disproportionate investment of the department in this area in the past. It is not obvious however that many issues that make sense in the context of logic programming are necessarily that significant in the context of other programming paradigms and many of the plans still involve logic programming in various forms.

Several logic programming research groups around the world have evolved toward constraints. This research group is not taking the same direction, on the basis that to do good work in this area, one must work on specific domains. One reviewer thinks that in fact the exact same observation is valid for abstract interpretation. Another possibility is to escape the first-order context, and to handle higher-order functions. The group has substantial experience in this area, exemplified by the work on lambda-prolog. However, this exciting direction of research has not attracted a very large following.

The research group led by Giorgio Levi has a reasonable size, and a very high international visibility in the logic programming community, and more generally in the field of semantics. There are many publications in excellent theoretical journals (Information and Computation, Journal of Logic Programming, Theoretical Computer Science, etc.) , and in high level conferences. The project has also international scientific collaborations, within Europe or with the United States, with researchers at the highest level.

The gist of the new direction is to bank on abstract interpretation as the common thread for all activities, and to rely on a substantial logic-programming heritage. Hence, the project is more geared to using and applying existing intellectual tools than to exploring new problems; it seems tool-driven rather than problem-driven.

As a result, it is not really clear that this positive effort to steer away from logic programming will really succeed. There should be some injection of ideological diversity,

maybe disruptive new lines of work. The new direction should also be formulated more concretely in terms of problems.

In truth, the ideas behind abstract interpretation are not new, and fairly intensely studied from a theoretical point of view. Many issues in program analysis, as was the case earlier in program optimization, revolve around efficient approximations, dealing with real programming languages, with floating point for example and not solely with integer numbers, with low level programming languages and with software (for embedded use for example) for which industry is actually willing to invest heavily in verification. This implies that the focus on strictly theoretical studies –often exhibiting intimidating abstractness– with little investment in pragmatic systems and actual software tools may not be the best research strategy for the future.

In this regard, a language like Java has a sufficiently reasonable semantic basis to make it attractive for study. There is a significant competition and even small progress on really important issues in practical analysis would be tremendous to advance the cause of the theoretical ideas that are proposed.

To sum up, the project comprises very smart people who have broad theoretical competencies. There is clear intellectual interaction with several projects in the department, and the desire to go beyond the earlier work on logic programming is to be highly commended. On the other hand, the new strategic direction that is proposed is not convincing and it would be preferable to refine the thinking on what problems need to be solved, before deciding which theory can be best applied to obtain the solution.

Project number 8. Algorithms and complexity: models and applications - Fabrizio Luccio

This is a major research effort, which has been ongoing for nearly a quarter of a century under the leadership of Fabrizio Luccio with the steady addition of competent collaborators, most of them former students in the same research group.

The group is well-connected with the international research community, publishes in the premier journals and conferences of the profession, such as the Journal of the ACM, the SIAM Journal on Computing, the Journal of Algorithms, Information and Computation, Algorithmica, FOCS, ACM Symposium on Computational Geometry, etc. The group is internationally respected as a high quality team, up to the best standards.

The main focus of the group is the design and analysis of algorithms, and the inclusion of "models" as a research item is more with reference to its algorithmic implications than as a self-standing topic. Algorithmic research remains a cornerstone of Computer Science research, despite the fact that technological progress still offers rapidly increasing resources. Even without the inevitable demise of Moore's Law, algorithmic design is crucial to the attainment of performance, and the work of this group assumes strategic relevance.

However, some critical comments may be raised in the constructive intent to further hone the quality. Although placed in the general area of algorithms and data structures, the research does not always offer the impression of being inspired by a "master" plan, directed at defining a field or a distinct topic in a field. In other words, one sometimes has the impression that the researchers essentially react to a number of interesting stimuli with no encompassing design. In addition, they do not seem to be sufficiently committed to the

current view that data structure research cannot proceed divorced from concrete applications (string B-trees being the exception).

One may also recommend that future members of the team be attracted from other institutions, with the aim of providing a diversity which is likely to further enhance the profile of the research.

Finally, one may note some yielding to the seduction of trendy topics. One such topic is DNA computing, which is insignificant from the computational viewpoint and is experimentally extremely critical: a clear assessment is in order before committal of any resource.

Different is the prognosis with respect to computational biology, and the group is encouraged -- in close collaborations with biologists -- to make appropriate investments (essentially by developing the necessary awareness of the research landscape) in order to make credible contributions to this critical field. Computational biology is a felicitous convergence of interests: While molecular biology cannot realize its goals without the partnership of Computer Science, it offers to computer scientists an unexplored wealth of problems, capable of fueling a new line of creative research.

Project number 9. Specification and Verification of Embedded Systems - Andrea Maggiolo-Schettini

This research project focuses on the use of temporal logics to verify properties of embedded software systems. There is a large amount of research in this area, especially if more broadly construed to include all reactive software systems (as suggested by Pnueli, Harel, and others) and to include verification techniques including model checking. The research group is led by three faculty members and has formed recently; we expect that the group will have to aggressively exploit their nascent connections with engineers and applications experts if they are to be successful in distinguishing themselves from the many other international groups competing in this area.

The group has strong credentials and backgrounds, having produced a number of theoretical results that have generally appeared in strong, internationally recognized conferences and journals. We have few comments on the specifics of the proposal. Our concerns are instead about the future of the effort: how will it have any impact beyond a very narrow research community?

The following comments should help in defining more precisely the specific contribution of the project:

- Why would it be easier to verify embedded systems than software in general? Since one view of an embedded system is a microprocessor with remarkably large amounts of loaded software, it is unclear why the "embedded" nature of the system makes a difference.

- The question of "where do the properties to check come from" is asked increasingly often now; the conventional answer that these properties are part of the specification is less and less satisfactory, since this is hardly ever the case in practice.

- As with other projects (such as #7, by Levi et al.), there is a key step of abstracting the important properties of a software system. If this is done a priori, then how can one ensure that that the abstractions are maintained and are useful later on? If it is done after the fact, then how can one ensure that the verification of these properties on the abstraction are true with respect to the actual system?

- How would this work interact with efforts in testing, in simulation, and in model checking, since undoubtedly actual software applications will be built using several, if not all, of these techniques? (Claiming that formal verification is the only "right" way to verify software is no longer acceptable, with the *possible* exception of small safety-critical systems such as nuclear power plant shutdown systems.)

- Despite the theoretical limitations of model checking, it may have great practical potential. (For example, the results of Chan et al. in Trans. on Soft. Eng. 24, 7 (July 1998) seem quite encouraging.) Even if this project does not pursue model checking, it should benefit from a detailed awareness of the current state of the art in that field.

Overall, we believe that the project pursues research on a very important topic, but one that is increasingly competitive and demanding. We recommend strongly that the researchers find and investigate some realistic applications and/or specifications; this is a significant effort on its own, but one that is guaranteed to have a large benefit in terms of the research that is produced.

Project number 10. The logic of molecule manipulation systems - Vincenzo Manca

This is a single-investigator project, which does not seem to involve a commitment of additional resources either from the department or from the public enterprise at large. Vincenzo Manca basically interacts with remote colleagues interested in similar investigations. To the extent that some amount of totally unfettered research is legitimate in an academic environment, this project fully qualifies in that category, and this opinion may be the end of our scrutiny. Vincenzo Manca is certainly well versed in the algebraic tools required by research in language theory, and no one can deny legitimacy to the pursuit of such investigations, as an expression of the mathematical component of Computer Science.

Language theory is a central component of theoretical Computer Science, and genuine connections with other domains, which would provide added relevance to the subject, should be pursued. However, there are well-founded reasons to believe that, in this case, the connections are somewhat artificial. This is an important issue within the scope of this evaluation, for it affects the credibility of the community at large with respect to fellow scientists in other disciplines. Indeed, this project is claimed to be relevant to biology and to be related to DNA computing. Both such connections are unsubstantiated in the proposal. The relation is metaphorical, since it is unconvincing that language theory has any relevance to the solution of genomic problems and even more tenuous is any connection to DNA computing: The latter is simply a cute instance of superminiaturization (fraught with enormous experimental difficulties) with no impact on computation models.

Vincenzo Manca should be encouraged to pursue his research in the honest setting of language theory. These remarks do not intend to infirm the validity of such research, but should be viewed as a negative recommendation for investment of resources in this area by the department.

Project number 11. Models and Languages for Open Distributed Systems - Ugo Montanari

This project concerns models, languages, and logics for distributed programming, particularly for large-scale distributed systems. It coherently addresses several related areas and problems. It includes foundational work on graph rewriting, coalgebras, higher-order, typed process calculi, and the "tile logic" model. It also includes work on the design and analysis of languages for distributed programming and on verification tools; some of this work has experimental aspects and generates some software artifacts (the KLAIM language, the JACK verification system).

This project is a continuation of activities that the same group of researchers has undertaken in the last few years. The group is of considerable weight, and one of Pisa's main sources of fame. It has international visibility in its area, with many publications in conferences and journals; it has developed a useful set of collaborations and contacts. Ugo Montanari is a prominent personality whose opinion and leadership have enormous impact in Pisa and beyond. He has substantial experience and a distinguished track record.

The proposal includes a broad, articulate description of problems in distributed computation, with a good discussion of WAN environments. While many of these are practical problems, the planned research has a solid theoretical foundation and also a theoretical orientation. Similarly, the frame of reference of the proposal includes a broad range of research projects but with an apparent emphasis on theoretical studies. The proposal outlines a principled approach, with remarkable mathematical sophistication, and with an excellent chance of yielding results.

The connections between the motivating problems and the proposed research are intriguing, but they often seem tenuous or indirect. This poses some dangers of over-abstraction and irrelevance; it also diminishes the likelihood of a significant impact on the world, and even on the theory world. The work on KLAIM, JACK, and other artifacts may help counteract these dangers; we expect that it will contribute in guiding the more abstract work and will include practical applications. Thinking about the actual problems of distributed systems, from time to time, may suggest useful research topics and opportunities for concrete impact. The project already has good plans in these directions, and we encourage it to push them further. They can make this strong research programme even stronger.

Project number 12. Advanced Methodologies Based on Soft-Computing to Develop Intelligent Systems - Antonina Starita, Alessandro Sperduti

This project is interested in combining symbolic techniques from AI and "Soft-computing", as represented by neural nets, evolutionary algorithms, genetic algorithms, and fuzzy systems. The subject area is addressed in two complementary ways: first in developing new ideas, for example how to use neural networks to classify structured rather than flat numerical data, second in developing applications where all such ideas are experimented with in the setting of an actual problem.

The project is of modest size, in terms of strictly academic staff, but its contracts allow for adding a number of people to participate in applications. There is a strong continuity between the planned work and the past achievements. The group is clearly aware of the field and it is publishing reasonably, for example in the IEEE Transactions on Neural networks, or at IJCAI, the premier conference in artificial intelligence. In comparison with

the teams that are quoted as reference for the field (for example Bishop at Microsoft Cambridge), it is probably not as strong technically as the best groups, in particular on the mathematics of the field.

The project is trying very seriously and competently to apply the techniques that it masters to applications in a variety of fields. One striking example is a system to analyze anomalies in the retina, in cooperation with the "Centro di Diabetologia" of the University of Pisa. Another example involves a cooperation with the "Dipartimento di Medievistica" for the inspection of paleographic documents. There is work on a number of other applications, showing a particular talent of this group for reaching out to other departments of the university, obtaining contracts from the European Community and cooperation with small industrial companies. This is viewed by all the evaluators as extremely positive for this project and for the department.

Of course, there are also possible difficulties, if the project should become involved in too many applications. First, it is necessary to accumulate competence, and if possible software, so that a basic library can be relied upon for each new application. This is apparently what is happening. Second, and this is somewhat more difficult, there should be some time available for a scientific appraisal of the methods and tools that are used. For example, in the retina application, success is claimed. But could success have been achieved as well by other methods, and how much has been gained by using neural nets in this application? In general, in this area as well as other areas of Computer Science, such as vision for example, there is a need for benchmarking experiments. This means that, for a variety of input data, one should pit one method against another and in this way obtain a more objective evaluation of their respective merits.

The Web offers an opportunity to test classifying and learning systems, and it is noted by the project proposal. On the other hand, working on medical applications, chemistry and data mining, while at the same time putting forward an original point of view on combining structure and soft-computing may be too much ambition for this project. A careful choice of applications, a focus on an original fundamental contribution and its impact on such applications should guide the strategic thinking of this project in the future.

Project number 13. ToMAKE: Tools and Methodologies for Data Analysis and Knowledge Extraction - Franco Turini, Dino Pedreschi

This project is a serious attempt to build on a long and distinguished tradition in the logic-programming area, moving the research into two important application areas: data mining and geographic information systems. We are encouraged by movement in this direction, and there are many positive indications, both in what has been done so far and in their future prospects.

The project has an air of solidity, with four faculty, a good group of doctoral students, and established links to several well-regarded logic-programming groups in Europe and the US. The two application areas --- data mining and geographic information systems --- are widely recognized as important database applications. We took note of the fact that in the mining area, which seems to have received the larger share of attention, there were two specific uses of their technology on behalf of a supermarket chain and the ministry of finance. The results in each case appear to have been satisfactory to the client, which is not always the case in data-mining endeavors.

A major strength of this project is the fact that it has already gone far along the road from the old style of theoretical research to a kind of computer-science research that is oriented toward applications and system building. We believe it offers a kind of training for students that properly mixes theory and practice. It helps the department build links to other disciplines that can use this technology, and it should be attractive to prospective faculty who have an interest in system building.

Once we consider this project to be engaged in applications, rather than a project about logic programming, the committee feels compelled to ask the natural question: is logic-programming the most effective way to solve the problems. It is true that there have been some recent developments, e.g., in LDL, that offer some powerful dictions for expressing the kinds of operations that seem useful in mining and GIS. However, what appears to be open is whether the implementation of these operations can be made competitively efficient. Both of the proposed applications are big-data applications, and efficiency needs to be considered (and measured) very carefully. Moreover, there are other technologies that have become available and deserve consideration. For example, there has been work under the VERSO project using constraint-solving to support GIS queries. The commercial relational-DB vendors are all promoting object-relational technology, which appears to make possible most, if not all, of the additional capabilities now found in LDL, while doing so in a framework that is likely to become a significant standard.

As a result of these considerations, there was a widely held sentiment of the committee that, while logic programming may have deserved a better fate, the adherence to this approach may not help the project achieve its goal of high-quality applications and tools. We hope that as the project progresses, it will be open to approaches of all types, according to the merit of the approach for a given application. In the best of situations, much of the technology that was developed in the logic-programming community will be transferred to other technologies.

The project offers opportunity for high-quality, innovative applications, and we give it our strong support. An over-dependence on logic programming is a risk. However, there is also an opportunity for the team, knowledgeable in logic-programming technology, to exploit that knowledge to develop software in more widespread technologies, e.g., object-relational database systems.

Project number 14. Architectures and Programming Tools for High Performance Computing - Marco Vanneschi

This is an umbrella project that includes a variety of efforts that focus on parallel programming and parallel architecture. It is a rather large effort, including five faculty members and a number of national and international collaborations. Parallel computing has, for many years, been an important research area, and this group has published many papers in the area.

The primary project in the parallel programming part of the proposal is the effort on "skeletons", which is intended to provide some idiomatic programming structures (roughly similar to design patterns) to solve common parallel programming problems. One example that was presented was the well-studied a priori algorithm, where the skeleton separated the completely parallelizable middle from the serial components at the beginning and the end. This example is a case where it may be worthwhile considering large serial machines (and memories) against parallel machines; even though this may not always be possible, comparisons against the serial case are often instructive, sometimes identifying additional

research problems. Other questions about skeletons include how one might handle inherently dynamic techniques, such as adaptive multigrid approaches.

Although the area of parallel computing continues to be important, the research landscape has been changing rapidly as the many advances in technology (e.g., the changing gap between processor, memory, and communication speeds) have come into play. Despite the many legitimate efforts in this project, we could not find a clear point of view about parallel computing throughout this proposal; this may be a costly absence given the continuous and significant changes in the underlying technology.

Furthermore, the most important long-term issues in parallel computation concern applications, not the underlying algorithms alone. We feel that is a weakness of the project that most of the numerical applications seem to be developed elsewhere than in Pisa. The proposal discusses at least one application, data mining, and hints at some others; we strongly encourage the group to consider these and other (perhaps scientific) applications in much greater depth, working to exploit the many groups that have associations with this project through its national and international collaborations. Even the notion of skeletons could be evaluated empirically more effectively by considering, for example, a very large numerical software application that has not yet been parallelized. Such a focus, combined with a clear-cut view of how changes in technology will affect the field, may allow this project to distinguish itself among the many international research efforts in parallel computing. Casting the ongoing research in the context of other international efforts would be useful both to the researchers in the project and also to others studying the results of the project.

One strength of the project is that its umbrella nature may provide a basis for work in building and evaluating systems. Indeed, when combined with a long-term vision of parallel computing and loose but careful direction, it may well be that low-level building and experimenting could lead to effective results.

***