EVALUATION OF INRIA-THEME ComB 2007

(Reseaux et Telecommunications)
Paris, 12-13 November, 2007

I. General Information

A. Introduction

INRIA has 132 research project teams (as of Nov 2005). The project teams are grouped into 16 programs which are in turns clustered in 5 major themes: (1) Communicating systems, (2) Cognitive systems, (3) Symbolic systems, (4) Numerical systems, and (5) Biological systems. INRIA maintains a unified view of the themes regardless of the location of its project teams. As part of INRIA's policy of research evaluation, the research project-teams are evaluated by an (external) evaluation panel every four years. This report gives the results of the evaluation of the Networks and Telecommunication Program (Program B) within the theme: Communicating Systems (abbreviated as ComB). ComB consists of nine research projects: ARES, DISTRIBCOM, HIPERCOM, MADYNES, MAESTRO, MASCOTTE, PLANETE, TREC, RAP.

In its analysis, the panel followed the general evaluation guidelines and criteria set by INRIA. The panel however felt completely free to express its views and findings in any way it wanted. The panel followed as systematic a policy as possible, to give equal attention to all projects. Two more projects belonging to the general networks and communications theme were examined during the meeting, but were not subject to evaluation (ARMOR, GANG).

B. The evaluation panel

- Sem Borst, Bell Labs, Lucent Technologies
- Costas Courcoubetis, CS Dept Athens University of Economics and Business
- Prof Serge Fdida, Université Pierre et Marie Curie - Paris 6, France
- Mario Gerla, UCLA Computer Science Department
- Piyush Gupta, Bell Laboratories, Alcatel-Lucent
- George Kesidis, The Pennsylvania State University
- Jean-Yves Le Boudec, EPFL Switzerland
- Alberto Marchetti Spaccamela, Sapienza Università di Roma
- Laurent Massoulie, Thomson Technology Paris Laboratory
- Sape Mullender, Center for Networks Research at Bell Labs
- David Wetherall, University of Washington , Intel Research Seattle

C. Theme ComB: Evaluated projects

- ARES - Architecture de Reseaux de Services
- DISTRIBCOM
- HIPERCOM - Communication a hautes Performances
- MADYNES - Management of dynamic networks and services
- MAESTRO
- MASCOTTE - Methodes Algorithmiques, Simulation, Combinatoire et Optimisation pour les Telecommunications
- PLANETE -Protocoles et Applications pour l'Internet
- TREC - Theorie des Reseaux et Communications
- RAP - Reseaux, Algorithmes et Probabilites

Two more projects - ARMOR and GANG -were examined by the Committee but were not subject to evaluation by the panel

D. The evaluation protocol

The evaluation process proceeded as follows. The evaluation panel was provided with activity reports for every project; the general assessment criteria of the theme; evaluation criteria for the project-teams, and information on INRIA's strategic plans for 2003- 2007. The panel also received a copy of the 2003 evaluation of ComB projects. On November 12-13 2007 a two-day evaluation seminar took place during which the panel met with representatives of all evaluated projects (and with several members of INRIA's executive committee and INRIA's evaluation board). The seminar consisted of:

- Day 1 (public sessions): a presentation of INRIA's overall research view of Com B area followed by half-hour presentations for each project summarizing the activities in the past four years and the future objectives, all in a plenary session with the complete panel in attendance.
- Day 2 (private sessions): 90 minutes meetings with (representatives of) each project separately, in each case with a pre-assigned group of three of the eleven panel members present. In the private session with a group of three evaluators, a project was given the opportunity to present a sample of its research tasks and achievements. In the same session, the group of three panel members discussed the project and its present and future course with the research team. The three evaluators of each project wrote individual summaries and sent these by email to one of them who was assigned as 'rapporteur'. The rapporteur synthesized the information into a report, which was reviewed and revised by the three evaluators and then transmitted to the coordinator of this report. The coordinator integrated the individual reports in the present overall ComB evaluation report, adding a commentary of the procedures used in the evaluation and extracting general observations and conclusion that cut across all projects. The full panel commented on the draft of the evaluation report before it was finalized. The inputs to the evaluators, including activity reports, synthesis forms, and slides of all presentations during the public and private sessions are accessible at the following website: http://www-sop.inria.fr/planete/evaluation2007/

E. Remarks regarding the evaluation protocol

We are grateful to the previous Panel of Experts in 2003 for requesting that each project report include a 'self-assessment'. In fact the 2003 Panel proceeded to supply a short list of indicative questions to be addressed by the project leaders. In the 2007 evaluation, the teams have complied with this request and have included corresponding sections in their reports. This Panel has found the sections on project positioning, collaborations, self-assessment and current positions of PhD graduates particularly helpful in its reviews.

In preparing its own evaluation reports, the Panel followed the evaluation criteria provided by INRIA, but expanded on them and divided them into a number of categories, so as to obtain a systematic and consistent framework. The criteria are similar to those used for the 1999 and 2003 evaluations and are listed in Appendix A.

II. Overall evaluation of the Com B Theme

In the overall evaluation of the Com B Theme we have chosen to follow an outline similar to that used by the 2003 Evaluation Panel. This will enable us to compare our impressions with those of the 2003 Panel. Specifically, we are interested in verifying that the comments and suggestions made by the last Panel have been implemented.

A. Vision, policy and leadership

The scope covered by the nine projects in Com B is very broad. The plenary introductory presentation by the INRIA Research Director did an excellent job in positioning the Com B theme in the context of the INRIA research mission. One issue is the correlation between Com B and other themes. Another important issue is the correlation among the projects within Com B itself. The fact that different projects are working on the same topics can be an opportunity for collaboration, or; it can lead to replications. A third option is that different projects are working on the same problem but with different approaches, thus leading to diverse solution strategies. Communications and networking techniques are the foundation of any distributed application. Thus, it is no surprise to see them appear in many different themes and projects dealing with distributed systems. At the top level, examining the main areas of research funded by INRIA, one discovers that Communications and Networking is covered also in the following

themes:

POPS: (Futurs/ComA) - System and Networking for Portable Objects Proved to be Safe, and:

RESO: (Rhones Alpes/NumB) - Protocols and software for very high performance Networks

In addition, one expects a close relationship between Com B and the projects of Theme C - Embedded systems and mobility

Likewise, within Com B, there is strong correlation among the 9 projects in many fundamental network research issues. Below is a sample of the correlations identified in the Plenary presentation:

- **Heterogeneous Networks** = ARES, TREC, PLANETE, HIPERCOM, MAESTRO, MASCOTTE
- **Overlay** = TREC, MASCOTTE
- **P2P** = MASCOTTE, PLANETE
- **Wireless** = MAESTRO, PLANETE, TREC, HIPERCOM, ARES, MASCOTTE
- **Ad-Hoc** = ARES, TREC
- **Mobility** = HIPERCOM, TREC
- **QoS** = HIPERCOM, DISTRIBCOM

The above topics reflect the most popular applications and related research problems in the field of Communications and Networking. The Panel finds that the nine ComB projects do in fact provide good coverage of all the key research facets in the theme area.

Further, the panel acknowledges that the projects in ComB have made significant contributions to the scientific community, often gaining international visibility and technical leadership in some of the areas. INRIA strategic goals have thus been met under the following criteria: coverage of relevant topics (as shown above); cross-team collaboration; collaboration with and transfer of results to industry; professional recognition team members, and; placing of its graduates in top institutions. More to the point, this Panel did not find any particular deficiencies in the coverage. The previous Panel, in 2003, critiqued that while ” *les systemes informatiques ubiquitaires*” (ubiquitous information systems) are mentioned as top priority in the Plan strategique, they are not references in any of the evaluated projects. This situation is now corrected, since at least 3 projects, ARES, PLANETE and HIPERCOM, deal such issues. Likewise, security, absent in the 2003 projects, is now well represented.

The 2003 report also commented that: ”*no new vision* was provided underlying the future course of the field. Rather *autonomously, similar research directions appeared in different projects* (like the interest for wireless networks) without connecting them”. This Panel thinks that there is merit in the ”*academic freedom*” research atmosphere fostered by INRIA, where each team is basically free, within general guidelines, to propose and pursue the research that they are most interested and best equipped to do. Since manpower is very limited, however, each team must carefully select a ”niche” area in order to succeed.

The 2003 Panel suggests that ” it will require much more ’Theme 1B-wide’ attention for relevant common aspects across project boundaries for INRIA to have an optimal and externally visible impact in current areas of technological advance and development”. The 2003 Panel continues to state: .. ”While there is clear advantage to *explore subjects like ad-hoc networking and Internet modeling with different methodologies* in different projects, there is a risk that hot subject areas are scattered over too many places. Integration efforts on common areas are encouraged, to radiate leadership in these fields and facilitate cooperation” Naturally, there is a tradeoff between research autonomy, that potentially might leads to fragmented peaks of excellence, and centralized management that seeks the integration of related expertises to achieve critical mass. This Panel recognizes that some projects have made major efforts to seek synergy and collaboration. It does, however, agree with the 2003 Panel that more should be done in this direction and encourages the teams to be proactive in exploiting the commonalities across projects. Specific comments/recommendations are deferred to individual project reviews.

B. The quality of research

One of the most accepted measures of research excellence is the number and quality of publications. From the data presented during the plenary Com B session, we learned that the total number of papers in 4 years was 296

journals (30% of which in major journals) and 838 conference papers (15% of which in very selective conferences with $<$ 20% acceptance rate). As a whole, this is an extremely good productivity for a total staff of 82 senior researchers, about one journal and two conference papers per year per person. Clearly, the productivity varies from project to project, with higher productivity in the more theoretical areas as we shall later discuss. Anther important measure (given INRIA education mission) is the PhD training. The projects support 82 PhD students from local Universities. Of these students upon graduation 75% move on to Academia and 25% to Industry. These are impressive statistics showing the critical role INRIA plays in Education and more importantly the formation of the French Academic force. Equally important is the fact that 50% of the graduates go abroad, helping maintain the visibility and image of INRIA in the world. At the same time, the researchers that opt to stay at INRIA often grow to positions of great prestige, such as elevation to the French Academy of Sciences.

Other elements that contribute to the quality of the program are the collaboration with Industry, the participation in European Projects - both successfully targeted by Project leaders. Associated to these activities is a substantial production of software systems and tools, and the participation in Industry Standard Committees.

C. Experiences with the evaluation protocol

We have carefully reviewed the comments and critiques of the 2003 Panel and are happy to report that proactive action was taken by INRIA management to correct many of the flaws identified by that Panel.

In particular, there was a complaint that the Plenary Presentations in Day 1 were "$\cdots$not focused on key issues; many projects did not articulate a good vision of what the project goals were". In the latest evaluation, all the presentations were well structured to convey up front the goals, position in the context of research in INRIA, France and worldwide, collaboration, funding, staff etc.

Naturally, each project leader felt compelled to report on the key "nuggets" of their research. In a half hour presentation with eleven projects to go, this can be very challenging. We sympathize with the comment of the previous Panel that "Four 'high-density' project presentations in a row without break easily turns a scientific experience into an endurance-test for a panel that wants to follow all presentations critically. A less intensively paced schedule would be more effective".. In fact, it would be good to limit the technical part of the presentation to one representative result that everyone can understand, leaving the rest of the technical results to the second Day discussion with a much smaller audience. The second Day, there is more interaction and the expert team can control the course and pace of the presentations to maximize the transfer of information

On Day 2, the time spent per project (90 minutes) is indeed $\cdots$ "insufficient for digging into the strategic motives and critical issues of a project (and in some cases, even discovering what they are)" $\cdots$ To remedy this problem, the 2003 Panel requested that $\cdots$"the synthesis forms (fiches d'evaluation) be not merely a presentation of performance facts and results, but also include a systematic selfevaluation" $\cdots$ This has actually happened. As we pointed out earlier, the assessment-oriented recommendations formulated in Appendix A of the 2003 Evaluation were followed by virtually all projects. This has helped the examiners to effectively zoom in on the relevant issues.

The 2003 Panel recommended that one or two selected publications be made available to the Examiners on Day 2. Apparently, this recommendation was not followed up, at least not by all projects. We think it will take minimal effort to add a couple of publication to the copy of the slides presented in Day 2. It will definitely help appreciate the style and quality of the research. Indeed, it should become mandatory for every project, as part of the evaluation protocol.

III. Evaluations of the Research Projects in ComB

In this section we give a detailed evaluation of the nine projects in ComB: ARES, DISTRIBCOM, HIPERCOM, MADYNES, MAESTRO, MASCOTTE, PLANETE, TREC, RAP Each evaluation loosely follows the criteria stated in Appendix B and ends with a number of recommendations.

A. ARES

Vision and Scope

ARES is a relatively new project that was established in January 2004. The scientific positioning of ARES is related to networked systems, with quite a broad focus, ranging from antennas to middleware. This strong, multilevel

system integration approach to address emerging hybrid and ad-hoc wireless networks makes ARES different from other Com B projects. On the other hand, this integrated vision is challenging and is a potential risk by itself.

The project develops its research vision through three complementary areas, (1) Wireless channels modeling, (2) Network architecture, (3) Services. The different goals are covered by addressing four priority targets:

- Integrating different types of mobility,
- Controlling cross-layer interaction,
- Providing self-configurability,
- Supporting quality of service (QoS).

The project supports experimentally driven research with many test-beds and prototyping.

Technical excellence

ARES is a fairly new group that is building up its research vision. It is composed of a group of young researchers with good reputation in the field and strong potential for growth. Its publication track is good though it lacks presence in major international conferences and journals. On the other hand, ARES is developing test-beds and software with a good impact and with potential for industrial transfers through industry cooperation. Its international visibility is growing and there are already signs of cooperation in that direction.

Importance/relevance

The different topics addressed are relevant and well suited to the vision driven by the scientific community. Some research activities are valuable; others lack inspiration or scientific methodology. The task on mobility management and integration is of interest and has to compete with a strong international on going effort in this area. The task on cross-layer integration is mainly devoted to wireless channel modeling and propagation simulation; it has delivered convincing results and operational tools. Its evolution towards MIMO should be better argued. The activity on self-configurability is more related to middleware and managed component-based services; though this work holds practical value, its scientific contribution is yet unclear. The group is particularly active and influent in the OSGI standardization. The contributions done by the last group, namely, network QoS and MAC protocols for multi-hop networks and wireless bandwidth estimation are of real interest and well-positioned at the international level.

Impact

The group holds a large set of industrial contracts at the National and European level. The team members are well positioned to go as far as proof of concept based on prototyping and software development. On one hand, the scope of the projects covered is perhaps too broad and carries the risk of lowering the critical mass necessary to address all the stated ARES priorities. On the other hand, the effort towards industrial cooperation and transfer is substantial and valuable. The publication record is in progress but should focus on higher-level quality conferences and journals.

Cooperations/coherence

ARES was created with the objective to develop an ”integrated” system approach to network design, ranging from channel modeling to middleware and services. This is an important vision; but, it is extremely challenging to develop, especially for a young team. ARES has developed a strong national visibility and is connected with many other labs at the national level. International visibility is building but still at an early stage. Industrial cooperations are numerous, both with national (France Telecom, Alcatel) or international (FP6 IST projects) stakeholders. The group is concerned to support a more formal approach of the problems through modeling (graph or queuing systems) but also cover experimental deployment and prototyping which is important in the research areas covered.

Strategic goals

The challenge of a system approach that involves multiple tasks and different skills has been identified as too difficult by the team members themselves and at risk with respect to critical mass. As a consequence, the project

leader has proposed to split the team into two sub-groups, each better focused on its own targets, namely, (1) ARES with focus on Wireless and Embedded Networking and (2) AMAZONES with focus on Services and Security.

Recommendations

The ARES project was founded in 2004 and build on a group of researchers covering a rather broad set of skills. It has achieved national visibility and some interesting contributions. The challenge of developing a system approach was found too hard and it was decided to split the project in two parts. The reviewers applaud this decision. The team members have a good potential. They should target an international visibility and clearly define their research agenda and priorities for the future.

B. DISTRIBCOM

1) Introduction: The following assessment was based on the DISTRIBCOM progress report dated Nov. 2007 prepared by the project's scientific leaders, Drs. A. Benveniste and C. Jard, and by the presentations given by the team during the INRIA COM B projects review held in Paris on Nov. 13 and 14, 2007. In the following, we make references to passages of their prepared progress report and references cited therein.

2) Progress Report and Presentations Assessments: Generally, the progress report is succinct but complete in that it identifies: key publications/contributions of the project in the context of the state-of-the-art (both methodological/theoretical and applications), collaborations (industry and academe), educational achievements (PhD graduates), and management structure.

The investigators have expertise in analysis by formal methods, control and modeling for distributed telecommunication systems. They continue to conduct basic research in these areas and document prior success in applications to digital communications and distributed alarm correlation, the latter in an optical networking context. These applications were described by C. Jard during the overall project overview on Nov. 13. The current research activity on web services composition, and document workflow for supply chain management was described in detail by A. Benveniste on Nov. 14. The reviewers were impressed with the more focused presentations on current application thrusts given on Nov. 14.

3) Specific Evaluation Criteria: **Adequacy to INRIA's scientific strategy**

The vision articulated is one in which the team will continue to pursue foundational research in their expertise domains of formal methods (DEDS, natural languages, Petri Nets, Markov chains, etc.). In addition, they are currently pursuing timely applications to web services composition and supply chain management.

Their core expertise is somewhat unique within INRIA. They have demonstrated the importance of their approach by successful application in several domains broadly interpretable as telecommunications and information technology. The applications sought are significant and timely in the opinion of the reviewers.

Generally, the reviewers felt that the DISTRIBCOM project's current and planned future efforts are very relevant to INRIA. In particular, potential outreach with groups in COM pursuing service/application-oriented networking could be very fruitful.

Scientific Excellence

Their foundational, methodological work is deep and is represented in recent PhD theses and publications. Their past applications to digital communications and signal processing received a best paper award from the IEEE Signal Processing Society ('03 transactions paper). Also, their mature work on distributed alarm correlation is timely and original and appeared in networking-themed conferences. Preliminary results of their current application thrusts have also appeared in leading software-related venues of IEEE and ACM.

Focusing on their current applied research thrusts, their work on document work-flows (docflows) in web/XML-based supply chains and orchestration of web services provides interesting, timely and important applications of, respectively, concurrent systems and hybrid systems (max/+ algebras, temporal ordering of QoS events, etc.). Web-services orchestration is based on service level agreements (SLAs), not explicitly on network resource consumption, with the goal of logically "safe" overbooking (i.e., efficient utilization of resources), and employing the Orc

specification framework (from UT Austin) to study practical examples. Other issues addressed by the project are attribution/identification, QoS assessments, sensitivity analysis, and reconfiguration policies in response to QoS failures detected on-line (i.e., on-line re-orchestration). This project will have impact, mainly do to their unique ”docflow” approach. The reviewers were impressed with both the practical nature of theses application thrusts and the sophistication of the theoretical tools being employed.

Industrial transfer and partnerships

They have a reasonable track record of outreach to industry and they are endeavoring to acquire more industrial partners interested in their current work. They plan to pursue open standards development with respect to their work on web services composition. Again, these reviewers feel that their application efforts are practically meaningful and are confident that their industrial outreach efforts will continue to be fruitful.

Manpower and means

The team seems to work effectively together and with several academics and industrial contacts (e.g., recently with UT Austin). They have significant outreach to researchers with similar expertise in Europe and North America. They are producing significant theoretical and applied results. Based on the presentations, the reviewers feel that the team possesses strong leadership. Again, their core expertise is complementary with other COM-B teams overlapping the subjects of distributed systems, formal methods/control, and telecommunication/information-technology. They have outreach to other groups within INRIA (e.g., the GEMO database group). The size of their research scope is consistent with the expertise and size of the team.

4) Recommendations and Concluding Remarks: Regarding their supply chain research (business workflow/docflow), outreach to companies such as Dell (via their UT Austin collaborators), and to local French companies in particular, should be further pursued.

The authors should continue to target publishing in ”applied” venues for their work on composed web services (e.g., those soliciting papers on software oriented networking, active networks, etc.). We also believe that fruitful collaborations are possible with researchers that are interested in related problems of network resources management for so-called service/application-oriented networks.

Generally, their pursuit of collaborations with ALU is encouraged and we note that they have had successful collaborations with industry in the past. The investigators further pursuit of basic research in their core areas of expertise is also encouraged.

Overall, the reviewers felt that both the basic research and applications to web services and supply chain management were very worthy and suited by the teams’ expertise and size.

C. HIPERCOM

Vision

As the project title tells, the goal of the Project is the study of High Performance Communications Networks. In particular, the Project has focused on Wireless, Mobile Networks. Mobility, scaling, radio propagation have been among the major challenges in this area, at least from the ”lower layers” perspective. The HIPERCOM Project Team has addressed all of these issues with a tri-pronged approach: (a) theory and design; (b) experimental verification and; (c) standardization and technology transfer. The strong emphasis on Standardization makes this Project unique among the rest of Comb-B projects. It is an important part of the ”vision” of the project leader. The acceptance of the network routing architecture OLSR by the IETF MANET group as one of the draft standards has contributed enormously to its popularity and to the international reputation of this Project.

Scope

The focal point of all the research activities within the HIPERCOM project has been the OLSR (Optimal Link State Routing) architecture. The architecture has evolved from one of the most popular wired Internet routing

protocols, the OSPF Link State Routing protocol. The HIPERCOM team deserves the credit for making important modifications to the ”wired” OSPF protocol, rendering it suitable for wireless and leading to a very efficient, scalable, robust routing protocol for ad hoc mobile networks. During the eight years of activities, the research has gone well beyond the design of a specific routing scheme and has addressed important, general issues in mobile networking, yet using the OLSR architecture as the testing ground for new ideas. The research has been based on solid theoretical grounds (graph theory, queueing systems, wave propagation, etc) and has led to a number of important theoretical contributions, including: reception models in random propagation fields, scaling laws for massive dense networks, asymptotic near-optimality of multicast reception with Network Coding. The theoretical results have all been applied (with no exceptions) to the OLSR model, leading to a number of new versions such as QOLSR (Quality of Service OLSR), MOLSR (Multicast OLSR) etc. In parallel, the HIPERCOM team has diligently lobbied the OLSR suite at the IETF MANET working group, and has influenced the shaping of important standards. The HIPERCOM team has paid much attention to the transfer of its technology to industry and military, with implementations ranging from battlefield to vehicular and airborne networks.

Scientific and technical excellence

In HIPERCOM, scientific excellence has been the foundation for a successful architecture design. The design of the OLSR suite has been based on the careful evaluation of alternatives using advanced performance evaluation techniques. In evaluating progress, the scientific goals set for HIPERCOM for the last four years were: Massive mobile dense wireless networks; New services and protocols; Wireless and backbone integration; Convergence of 4G and MANET Internet.

There was strong theoretical progress in the first two areas, leading to extensions of the Gupta-Kumar law to non-uniform traffic distributions and to multicast capacity. In the third area, the results are of a more specific and applied nature, leading to extensions of the OLSR architecture such as Quality of Service and Security in OLSR. Contributions to IETF such as autoconfiguration of OLSR networks are also noteworthy. We did not see reports on the convergence of 4G and MANETs. This is definitely a hot topic, however in our opinion it is currently more of political than research relevance. When the 4G architecture stabilizes, it will make sense to study convergence. Thus, it looks like a topic to be carried further to the next phase.

These results have been reported in an impressive number of publications (over 30 journal and over 130 Conference papers over four years). Further dissemination of these results has been accomplished through lectures at leading Universities in the Country. Presence at IETF has enormously increased the visibility of the team. The overall reputation of this effort is proven by the more than 500,000 hits on Google. One must also stress the fact that this team has managed to maintain a good balance between scientific excellence (making fundamental contributions to the field) and technical leadership (by delivering real products to the community).

The plan for the next four years calls for: Theoretical limits of mobile networking; OLSR and related protocols; Beyond OLSR and Internet; Drone and vehicular communication.

This is an excellent plan, which again strikes a good balance between theory and practice. However, the Reviewers believe that the HIPERCOM Team has already proven itself in the areas of routing and ad hoc network layer theory. They recommend that the Team, while protecting the OLSR investment, take a broader, more aggressive look at the new, exciting research issues that are now emerging and are impacting the design of future large scale ”opportunistic ad hoc networks” - from Physical/Network layer cross layer design to emerging P2P mobile applications. Probing these areas will provide further opportunities for collaboration with other Com-B teams such ARES for the Physical Layer and PLANETE for P2P.

Adequacy to INRIA’s scientific strategy

The HIPERCOM team has made very significant contributions to the Networking Field: they have an enviable publications record; a history of OLSR adoptions in military, civilian and academic institutions; a very strong international reputation; an exceptional ”citation” record. These accomplishments are all in perfect accord with INRIA scientific goals. What is missing? This Review team believes that after 8 years of OLSR success, the HIPERCOM researchers are ready for new challenges. They must expand the scope ”beyond OLSR” to the next ”killer” architectures and applications. In this pursuit, they should seek stronger collaboration with other INRIA

teams (eg ARES, PLANET etc).

Manpower and means

The international visibility and reputation of the HIPERCOM Team is strongly tied to the Presence at IETF, to the success of the numerous OLSR implementations all over the world and to the large number of ”followers” who are extending OLSR in new directions. Thus, we have taken note of an excellent presence in scientific forums. The academic image of the Team could be improved - leveraging/expanding the University contacts and increasing collaboration with academic researchers in France and abroad. To that effect, the presence of Prof Thomas Clausen, Ecole Polytechnique, and Prof Khaldoun Al Agha, University of Paris, are important assets to the Project.

Recommendations and Concluding Remarks

To conclude, the Review team is impressed with the accomplishments of the HIPERCOM Team. It recognizes the fact that the goals proposed four years ago have been by and large achieved. The success of the OLSR architecture has surpassed all expectations, with numerous technology transfers and adoptions in diverse applications domains. The Team has not ”rested on the success of the 500,000 Google hits” but has proactively explored new theoretical grounds in large scale networking, systematically utilizing the results from theory to build a more powerful and versatile OLSR architecture. As a whole, the HIPERCOM activity has enhanced INRIA’s visibility in the scientific world and has helped INRIA achieve its strategic goals.

The Review team believes that the HIPERCOM project will continue to prove successful and productive over the next 4-year cycle. However, the Reviewers suggest that the Team, while protecting the OLSR investments, take a broader view and consider the new technologies that are emerging in the mobile network field and that are likely to have major impact on future large scale mobile network designs - from cross layer optimization (eg MIMO and Network layer) to emerging P2P mobile applications. Probing these areas will provide further opportunities for collaboration with other Com-B teams. It will also be essential to maintain the Competitive Edge in the next four years.

D. MADYNES

Vision

The MADYNES project started in January of 2003. Its vision is research of management and security in networks and networked services. A main goal was to develop new techniques for dealing with scale and dynamics in large networks. The project does this by (1) developing techniques to allow nodes to monitor, manage and fix problems by themselves, (2) techniques for managing and assessing security in large-scale systems and (3) designing key distribution and revocation protocols for multiparty services at large scales. To the extent possible, the project seeks experimental validation in wireless and wired networks, IPv6 and VoIP.

Scope

MADYNES represents a significant contribution to INRIA strategy. The problems addressed and approaches taken distinguish the project from other projects enough that there is little concern that MADYNES duplicates work in other projects. The project addresses four subfields: (1) Disruptive Management Approaches, (2) Integrated Management, (3) Key Management for Dynamic Communities, and (4) VoIP Security. The problems addressed in MADYNES are important ones and the approach used by the group is eminently practical. This is borne out by the group’s publication record. In the self-assessment report issued by the MADYNES team, there is a good overview of related research and competing research groups.

Technical excellence

(a) **Disruptive Management Approaches**: It is not clear to the evaluation team what is meant by ”disruptive” management. Presumably it has something to do with managing nodes in a peer-to-peer network and thus with

choosing a manager among peers. What management then takes place is left to the reader to speculate. With the information provided by the project team, it is hard to evaluate the work or its significance. Furthermore, the evaluation team feels that a large body of literature on electing ”leaders” among dynamically changing groups of nodes is being ignored. (Chapter 9 of Elements of Distributed Computing by Vijay K. Garg (Wiley 2002) discusses many.)

(b) **Integrated Management** The approach in this subproject, whose title does not quite reflect its content, is to create a method by which failing nodes in an ad-hoc wireless network can be discovered. The problem is to distinguish nodes that are merely out of reach of a forwarding node from those that misbehave, e.g., because of bugs, or battery failure. The project team proposed using a Markov chain model to characterize ”pathological’ nodes. The evaluation team could not discover any reason why this was a reasonable choice, and this is cause for concern. The evaluation team sees two choices for the project team: either abandon the method altogether or come up with a convincing model for the distinction between failed and out-of-reach nodes.

(c) **Key Management for Dynamic Communities** In this subproject, the goal was to find algorithms and methods for distributing group-communication keys. The evaluation committee found several points where clarification would be desirable:

- It isn’t clear whether mobile ad-hoc networks are sufficiently different from other networks to warrant developing a key management protocol targeted specifically to these networks.
- It is not clear whether the solution presented is the best one and we saw no arguments defending the particular choices made.
- The initial protocol had errors that a security-verification tool uncovered. It is not clear if there are other, undiscovered errors.
- The example scenario presented was that of content distribution, for example streaming audio or video. It is no clear whether revocation is required in this particular scenario: One could imagine each song (or number or item) being encrypted with a different key-a key that is distributed only to those nodes that have access to the song. Then there would no longer be any need for key revocation, nor would there be a need for ”subkeys” used in smaller groups.

The evaluation team felt there was still a lot to be done in this project. Given that this project has been running since 2003, the team was somewhat disappointed both in a qualitative and quantitative sense.

(d) **VoIP Security** In contrast to the title, this subproject is primarily about functional testing of SIP implementations. The project team came up with a combination of methods-brute-force, fault-injection, fuzzing-with which commercial Voice-over-IP systems have been confronted. Serious implementation errors were found in many of these systems. All in all, the team came up with more than 30 successful attacks. The subproject deserves to be the flagship of Madynes. It has given the team significant international visibility as well as acknowledgements from a number of VoIP equipment vendors. The project claims the first ”stateful” fuzzer for SIP. This allows far deeper probes into implementation weaknesses than competing attempts. The team has also built an application that constructs (parts of) the state machine that drives the SIP protocol in the system under investigation. This gives an interesting handle on future probes into different systems.

Importance/relevance

The biggest success story of the Madynes project is the work on SIP fuzzing. This has led to significant international recognition, in which the project team is clearly recognized as the leader in the field. In the other subgroups the work has become much less visible. Most of this work can be characterized as experimental and does not exhibit significant interplay between theory and practice.

Impact

The group as a whole has a good publication record with roughly one journal publication per senior member of staff every two years and two conference papers per year. The group has also produced a respectable software portfolio, much of which is made available to the general public. The evaluation team recognizes the effort it takes to build, distribute and maintain software distributions. The group has been active in the research community by

participating in reviewing and conference organizations. By far the most striking result from this group has been the work on SIP vulnerabilities. This work has brought the team to significant international recognition. The other subprojects within Madynes, unfortunately, remain much less visible.

Cooperations/coherence

Madynes is very well connected internationally, with a strong role in the community as illustrated by the leadership of the Emanics NoE and a strong involvement in IFIP WG6.6. The team members are familiar with the important researchers in the field and they collaborate with them where appropriate. There appears to be no important reason why the four subprojects in Madynes belong in one umbrella project, except for overlap in the research staff. We do not see this as problem.

Strategic goals

The plans for future research have been phrased mostly in the context of management. A number of different management approaches are being proposed: managing uncertain environments, self management, botnet tracking for large-scale management, etc. But it is left unclear what management consists of. This detracts from the value of the proposed work in that it becomes hard to measure success when the goals are so unclear.

Recommendations

The Madynes project has important strengths as well as weaknesses. The work on fuzzing is well recognized internationally and needs little advice from the evaluation team, except perhaps the warning to the project team not to rest on its laurels. For example, it could use the fuzzing approach in other protocols than SIP and reduce the human effort in finding the weaknesses in such protocols (i.e., to carry the automation of the fuzzing process further). The other three subprojects, however, must improve. They cannot continue as they are. The evaluation team recommends the following:

(a) Disruptive Management Approaches: the goals must be stated much more clearly, as well as the criteria for what constitute successful management algorithms.
(b) Integrated Management: Key to this project is distinguishing between ”normal” and ”pathological” nodes. Merely using a Markov chain model will not do. Either better distinguishing features need to be found or the approach should be abandoned.
(c) Key Management for Dynamic Communities: this subproject needs two things: a motivation for the need for specialized key distribution protocols in ad-hoc wireless networks and a far more rigorous approach to defining and analyzing the protocols.

Further remarks

This project is all about management, and management is a fuzzy concept. For example, one can say that ACK packets ”manage” a TCP connection, but TCP ACKs are never considered management. Andrew Herbert once defined management as ”all the things we don’t have an algorithm for” and that is still a useful definition. We think this project and many other research projects like it could greatly benefit by the use of a more carefully chosen term than ”management”. Terms like ”distributed failure detection” or, ”performance measurement” come to mind. As Olivier Festor remarked during his presentation, ”Network monitoring and defense approaches” covers the intent of the project much better than ”management and security”. Most of the work presented to us concerns measuring various parameters in networked systems; little of the work involves acting on those measurements.

E. MAESTRO

Scientific Excellence

By its scientific production, this team ranks among the very best worldwide in its area of expertise. This can be attested from the quality of the publications, at top venues in the networking domain such as ACM Sigmetrics and IEEE Infocom, as well as top theoretical journals, such as Internet Mathematics. The scientific excellence and visibility within the international community is further testified by the strong representation on editorial boards of flagship journals and numerous program committees of major conferences. The range of topics addressed by the team over the last four years is also impressive, ranging from transport control design to 802.11 modeling and performance analysis, game theory, and web page ranking algorithms among others. This team provides a unique blend of theory and practice, which makes their results extremely impressive. The breadth of its theoretical underpinnings is quite substantial (ranging from classical stochastic process theory to Markov decision processes and evolutionary game theory). It is quite remarkable that they could apply many theoretical tools to relevant, hot problems.

Adequacy to INRIA's scientific strategy

Members of the group have been extremely agile in their adaptation to new topics, as can be seen from their recent interest in web page ranking, content placement in CDN's, and content delivery in "delay-tolerant networks". The long-term research vision of the group is well articulated and thought through, and their general focus fits perfectly with the general strategic objective focused on networks and communications. Maybe some more emphasis on security related performance issues could be encouraged. Their recent effort within the framework of the "Bionets" European project is also in line with the strategic objective around modeling for control of complex systems. The team's contribution to Bionets stands out by its depth and excellence and this direction seems quite promising.

Industrial transfer and partnerships

There have been several formal collaborations with France Telecom R&D, which led to the design of scheduling strategies for prioritizing file transfers over communication bottlenecks such as DSL access lines. While further collaborations with FT R&D are likely to take place, the form this will take is yet to be defined. There is also a new collaboration between A. Jean-Marie and a start-up, concerning content placement for emerging Content Delivery Networks. A "virtual" INRIA-Alcatel lab is being launched. The MAESTRO team plans to do work on wireless technologies, and more specifically mesh networks as well as self-organisation for network control as a contribution to this joint lab. This constitutes a very substantial industrial collaboration effort, for a team of this size, and whose scientific contributions span all the way from the applied to the more theoretical aspects. Besides the extensive industrial partnerships, the group is also actively engaged in a wide range of collaborative efforts with other leading research groups in Europe and around the world.

Manpower and means

The permanent members of the team have done excellent individual as well as joint work. It would seem they have the "critical mass" and internal coherence for success, which is a particularly noteworthy achievement, given that they are scattered over three different sites.

F. MASCOTTE

Vision

MASCOTTE's vision is to utilize tools from discrete mathematics (in particular, graph theory), algorithms, combinatorial optimization and simulation to address issues in telecommunication and transportation networks.

Scope

Mascotte does not design protocols but aims at constructing network and communication algorithms. In the last few years the project has developed both theoretical and applied tools for the design of heterogeneous networks

consisting of various types of communication systems, eg, WDM, SDH, wireless, satellites, etc. The project also addresses overlay networks and fundamental issues in discrete mathematics and algorithm design. The areas addressed by the project are excellent niche within INRIA with no or little overlap with other projects, and of value to advancing the state-of-the-art in network optimization and networking theory.

Technical excellence

MASCOTTE team members are well-known in their core areas (graph-theory, combinatorial optimization, algorithm analysis etc.). The group has maintained a high publication record in both top and very good journals related to the core mathematical areas of the Group. The senior group members are internationally well-known researchers. The group has a good amount of collaborations within France and with industry. The team's publications (particularly the quality ones) are concentrated in the more theoretical areas; in particular the team has many publications in flagship journals such Combinatorica, Journal of the ACM, J. of Combinatorial Theory, J. of Algorithms, ACM Transactions TOPLAS, Information and Computation. In the more applied areas, the team has publications in flagship networking conferences such Globecom and IEEE Infocom and the overall quality is very good.

Importance/relevance

The topics addressed in backbone networks, wireless networks and fault tolerance are relevant from a research perspective, and some of them also are very relevant to the networking industry (traffic grooming, WDM reconfiguration, wavelength allocation in optical networks, energy awareness in wireless networks). The simulation work benefits from the strong theoretical background of many of the project members in supporting more applied work done in cooperation with industry. Though good, the integration of the simulation work with the rest of the project could be improved. The activity on overlay network is a hot topic, which started within the project in a prototypical way and will require more resources to grow.

Impact

The group as a whole has a very good publication record with more than 60 publications in good journals, and many flagship publications in the more theoretical community. The several industry collaborations on network design and optimization are indicative of the good impact in the industrial application of the mathematical work. The group has also produced some software, most of it in cooperation with industry. The group has been very active in the research community by participating in reviewing and conference organizations.

Cooperations/coherence

The project team members have good academic research collaborations and also good industrial collaborations. The impact on networking would be much better if the project were to have stronger linkage with system-oriented networking groups within INRIA and/or other academic institutions. The overall project coherence and impact would benefit by focusing the research activity on fewer themes within each objective.

Strategic goals

The overall strategic goal is to continue addressing networking problems with theoretical tools and proposing solutions that are validated through software implementation. The proposed plan is coherent with the very good work carried in the last few years. However, some of the specific objectives for the future can be improved by having more networking focus. In fact some of the proposed items require a good knowledge of constraints and issues that should be included in the modeled problems; otherwise, the risk is to study theoretical problems with limited impact in the network community. Closer cooperation with systems-oriented networking groups might be helpful in increasing the impact of the considerable mathematical skills of the project team members. The work carried on overlay networks is important but is less coherent with the planned work in other topics. Given also the limited number of researchers currently involved in the activity, the need of cooperation with other researchers is

particularly significant for this activity.

Recommendations

The project is of very high-quality in the mathematical areas, and its impact would benefit by concentrating on fewer areas driven by important problems in networking research. In particular, the activity on overlay networks could better accomplish INRIA's strategic goals by strengthening connections with other researchers. Overall, the project members will benefit from closer interactions with systems-oriented research groups both within INRIA and outside.

G. PLANETE

Vision and Scope

The main aim of the PLANETE project team is to propose and study new architectures, protocols and services that will enable efficient and secure communication in heterogeneous networking environments. The primary emphasis is on the Internet. The PLANETE project was started in 2001, growing out of activity by the earlier RODEO team. It was previously evaluated in 2003. The team is medium to large in size (10 PhDs and 4 engineers plus students) and spans INRIA Sophia Antipolis - Mediterranee and INRIA Grenoble - Rhone-Alpes. PLANETE has a strong reputation, owing to its strong roots and continued productivity. It emphasizes the systems side of communications research, including the development of experimental prototypes and software toolkits. It participates in standards bodies such as the IETF as a means of obtaining impact. It differs in these respects from most of the Com B project teams, with the notable exception of HIPERCOM and its OLSR efforts. This is a valuable addition of diversity to INRIA's Theme Com B research portfolio. Further, PLANETE is the only Com B team working on the topics of wireless security and file delivery in broadcasting systems. PLANETE competes with experimental network research groups in universities and research labs around the world, with the specific groups varying with the specific topics of study.

Technical excellence

Over the study period, the PLANETE team has delivered many individual results of high quality in the sub-areas in which it has concentrated research activity:

(a) Security in infrastructure-less networks. Research examined issues of trust in RFID and wireless sensor networks as domains, as well as MANETs. This resulted in the development of novel techniques for secure data aggregation in sensor networks, and secure device pairing.
(b) New modes of information dissemination. This includes digital video broadcasting and peer-to-peer techniques for content replication. Notable here are the FLUTE application for reliable file delivery, as well as component technologies such as the development of efficient LDPC erasure codes. Efficient self-verification of these codes is an innovative development.
(c) Multimedia transmission over heterogeneous networks. This includes work on scalable overlay techniques as well as cross-layer (MAC/PHY) optimization. The work in this sub-area is good, yet there are few highlights to call out.
(d) Understanding Internet behavior. The focus is on scalable measurement-based monitoring methods and measurement-based services for applications. Notable here is the development of secure Internet coordinate systems and research on understanding the incentives of the Bit Torrent peer-to-peer system.
(e) Networking evaluation platforms. This includes experimental testbeds and realistic simulations. Notable here are both the OneLab extension of PlanetLab and significant contributions to the ns3 simulator.

This technical excellence is evidenced both by scientific publications, and by the collaborations and industry endorsement of research outcomes (as described under Impact below). During the evaluation seminar, the reviewers had the opportunity to interact with part of the PLANETE team to examine its work on three specific research topics: securing Internet coordinate systems; wireless sensor network security; and clustering and sharing incentives in Bit Torrent (the most popular peer-to-peer system). This examination left the reviewers confident of the technical

excellence of the work. In each case, the reviewers were impressed by the competence of the presenters on their technical topics and handling of questions, including their knowledge of competing approaches and other recent developments in the state-of-the-art. The reviewers were also impressed by the collaborations with top external research groups, specifically UCLA for the Bit Torrent work and UCI for security. The team is producing successful PhDs who have obtained positions at NICTA (a leading Australian research institute for IT), EPFL, University of Michigan, University of Liege, and Thomson.

Importance/relevance

Much of the PLANETE activity is leading edge work that is addressing what the research community believes to be the important problems. Research on OneLab, ns3 and involvement with FIRE (the European counterpart to GENI and initiative on Future Internet Research and Experimentation) is focused on the state-of-the-art in experimental research platforms. Research in the wireless area is focused on the problems of security and includes important technologies such as RFID. This is an important problem area, and it also provides a good counterpoint to much other research on the performance of MANETs. Research in the area of Internet measurement has been applied to the development of application services (a secure coordinate system) and the analysis of peer-to-peer systems. These are both important problem directions. The former, the application of measurements to services, is a direction that is likely to be of increased importance as Internet measurement moves beyond a tool for study. The latter is a valuable as a study of the single largest contributor to Internet traffic today. Other research is focused on content delivery - reliable multicast protocols, applications, and supporting technologies such as FEC codes. Content delivery is and will remain an important topic. However, the PLANETE team should not limit itself to reliable multicast as a distribution mechanism.

Impact

The team has achieved a high level of impact through strong ties to industry and standardization bodies, and through scientific publications in highly regarded venues. These provide it with visibility at the international level. The team published 31 journal articles and 84 conference papers over the evaluation period. Its work appeared in premier venues for the work at hand including SIGCOMM, IMC, SIGMETRICS, JSAC, TON, INFOCOM, and Mobisys. Shifting more publications into top-tier venues, especially at experimental network research conferences (SIGCOMM, NSDI, CoNext), would further boost the visibility of the team. Many of its collaborations are top notch, e.g., Princeton for the extension of PlanetLab, UCI for cryptography, UCLA for Internet measurement, the University of Washington for ns3. Much of its research output is either picked up by companies, e.g., the open source LDPC codec that is not encumbered by IP, picked up by standardization efforts, e.g., the FLUTE file delivery system by IETF and DVB, or picked up by the research community, e.g., ns3.

Cooperations/coherence

PLANETE appears to be well managed, especially given that it is a medium to large size project team that is split across two geographically separate locations. An issue for the team is the relatively large number of research thrusts that are ongoing. The team is well-aware of this issue and has been consolidating its research focus over time. It uses mechanism such as regular audio conferences to keep all team members in touch with each other. We would recommend that this continue. As a group, the PLANETE team is aware of the topic overlaps with other INRIA groups and appears well informed of the research activities of those groups. There has been little research leverage between PLANETE and other groups to date though (at least of direct exchanges of which the reviewers are aware). Externally, PLANETE is very well connected. It has top notch collaborators, many of which are international. It is an active and involved participant in international standardization groups, such as the IETF, IEEE (for 802.11) and DVB. It appears to have used ties to companies well. PLANETE also has visibility as a group in the broader research community through its release and distribution of research artifacts.

Strategic goals

The PLANETE team proposes to target two high-level goals looking ahead: the development of portions of the future Internet architecture, focused on secure communication and efficient data dissemination; and the development of network evaluation platforms, including experimental platforms and simulation tools. These goals support each other, as the second may be used in support of the first. The team further proposes to focus its research in support of these goals in four sub-areas that grow out of its research competencies to date:

- Data-centric networking
- Network security, especially in wireless networks
- Network monitoring
- Evaluation platforms and methodology

The reviewers considered this to be a sound plan. Each of the topic areas is focused on an important problem in its own right. The team has relevant expertise and competence for each topic. And the proposed research is complementary to other Com B research activity yet contributes well to the larger INRIA strategic objectives. Moreover, this activity plan pushes in the direction of consolidating research activity around a smaller number of high-level goals and sub-areas. This is a continuation of an ongoing effort to build more focused research interactions across pieces of the larger project. This is a direction that is recommended by the reviewers as a means of improving the cohesiveness of the project and increasing its likely impact.

Exploring the topics in slightly more detail, data-centric networking and network security are both key issues for the future Internet architecture, and research on them is timely. Research on evaluation platforms and methodology for using them effectively is an important component of making progress on these issues. Again, it is timely, given the emergence of GENI and matching European initiatives. Research on network monitoring, however, sits between these categories as neither a problem nor a development tool. Thus we suggest that the PLANETE team consider how to apply network monitoring in support of their other goals, e.g., to security or data-centric networking or as a component of evaluation platforms. This will help with research focus.

Of the planned research mentioned by PLANETE, the part that seems most difficult to fit into the two high-level goals is that on RFID and sensor networks, as the objectives are stated around future Internet architecture and evaluation platforms. This is worthy research, and the reviewers do not suggest that it be discontinued, but rather that the team looks for other integration opportunities. For example, one good reason to include this wireless research is to emphasize the future Internet as ”an Internet of things.” This would suggest the consideration of how RFID and sensor systems interface with the rest of the Internet rather than operate standalone. Similarly, data-centric networking is a thrust that could fit well with these topics, as could experimental platforms and simulation. These are only suggestions; the team understands the issue and must devise a fitting response.

We also note that the proposed research directions are conservative in the sense that they build directly on the competencies of the research group as developed in its research to date. That is, they are unsurprising and a linear extrapolation of where the team has been, other than the new thrust on network evaluation platforms. We do not have recommended changes in mind, as the set of topics is fitting. Nonetheless, we would encourage the team to deliberately decide at a more detailed level within topics what new lines of work to start, and what existing lines of work to stop, lest the decision be made for them by default. For instance, work on data-centric networking could continue to build on a large body of work by the group in the area of reliable multicast. Or the work on reliable multicast could be declared to have succeeded, freeing research on data-centric networking to take new directions. The latter would introduce more risk and more opportunity.

Recommendations

To conclude, the reviewers find that the PLANETE project has delivered a strong performance as a research team over the evaluation period and is functioning well as a group. The experimental network systems thrust of the group is valuable part of INRIA’s Com B research program, adding significant diversity to the portfolio. Based on the strategic goals proposed by the team, the reviewers further believe that the PLANETE project will continue to prove successful and productive if it is allowed to continue for another four year cycle.

This report is largely in harmony with the self-assessment of the team (including its future plans), as delivered in its written report and presentations during the evaluation seminar. This harmony does not come from a lack of opinion on the part of the reviewers, but mainly comes because the team is both very successful and has openly received the feedback obtained at the last evaluation cycle and used it to further strengthen its goals and operation.

Thus the main issue that the reviewers wish to highlight as part of the overall recommendations is one of which the team is already well aware – breadth and the existence of many diverse research thrusts.

The research of the PLANETE team is spread across several topics and is broad, with research in the group on wireless security for RFID and sensor networks seemingly distant from much other research also in the group such as Internet measurement and the investigation of incentives in peer-to-peer systems. It is important that the team pursue a small and cohesive set of goals: this will be the means by which the team can take its impact and visibility to the next level. The PLANETE team appears to realize that they stand to gain from consolidating their research activity, and they have made progress in this direction over the review period. They plan further consolidation looking forward (see Strategic Goals above) in terms of the number of topics addressed, and in the establishment of two overriding objectives for the project. The reviewers considered this to be a sound plan and encourage the team to look for further points of interconnection, e.g., as occurred by targeting the use of Internet measurements towards peer-to-peer services.

The reviewers felt that the team should continue to target top publication venues to achieve impact, as well as develop and release research artifacts infrastructure and participate in standardization efforts. The publications of the last cycle (including venues such as SIGCOMM, IMC, INFOCOM, CoNext, SIGMETRICS, JSAC and TON) have served the team well. Shifting more of its publications into top-tier venues to improve the publication record further, especially at experimental network research conferences (SIGCOMM, NSDI, CoNext), will boost the visibility of the team.

Further Remarks

As part of the evaluation seminar, the team was forthright in discussing the recommendations and concerns that were aired in the previous evaluation cycle and how they had responded to them. The reviewers found that PLANETE team has done much to address previous feedback directly (several points are noted above) and commend them for it. This effort extended into the evaluation seminar, with a well-organized and written self-assessment report and well-prepared presentations.

H. RAP

Vision

The RAP project started in 2004, with the primary charter to foster and strengthen the fruitful collaborations with France Telecom R&D, in particular Fabrice Guillemin and colleagues. The main thrust is the analytical modeling of communication networks, with the two-fold purpose of evaluating performance and devising algorithms for controlling, designing and operating such networks. The project plays a highly valuable role in capitalizing on the mutual interest of the France Telecom and INRIA team members, and bringing together their specific, complementary expertise in operational aspects and advanced probabilistic methods, respectively.

Scope

The project targets highly relevant problems, both from a theoretical viewpoint and from an application-oriented perspective, and has delivered excellent scientific results. By its very nature, the project makes a strong contribution to the INRIA strategy, focusing on fundamental research contributions, and leveraging these to address difficult problems of significant practical relevance. The main theoretical thread concerns the advance of scaling laws for stochastic networks, while the more application-oriented goal pertains to the analysis and design of traffic measurement algorithms in high-speed commercial IP networks. Further important activities that have been pursued include the analysis of bandwidth-sharing policies in IP networks with streaming sessions and elastic applications, and the analysis of so-called tree algorithms. The proposed future research directions seem a natural extrapolation of the ones for the previous four-year period, placing less emphasis on bandwidth-sharing in IP networks and tree algorithms, and stepping up the effort on resource allocation mechanisms, with network security as an important and mostly new line of research. Some of the topics addressed by the RAP project are closely related to research pursued by other INRIA projects, particularly MAESTRO and TREC. The team members are well aware of the

activities of these other projects, and past experience (e.g. in the context of TCP modeling) suggests there are strong benefits to addressing the same subjects from different angles and using different methodologies, so duplication is of little or no concern.

Technical excellence

The scientific accomplishments of the RAP project are of high caliber, in terms of technical quality, depth, and originality. The results have been published in renowned discipline-based journals such as Annals of Applied Probability and Random Structures and Algorithms, as well as some of the major application-oriented conferences in the area, such as Mobicom and Performance. The ability to operate so successfully at the forefront of both the theoretical research field and in the engineering/networking arena is most impressive, and achieved by few groups in the world.

Importance/relevance

Some of the most remarkable results include the scaling limits for wireless cellular networks with user mobility and call blocking, which revealed some intriguing multi-stability phenomena in a new context, distinct from the traditional setting of dynamic routing in loss networks. While mostly theoretical in nature, the qualitative insights and practical implications have also captured considerable interest from an application-oriented audience, as illustrated by the Mobicom paper. The research activity on tree algorithms has yielded further striking results, providing an entirely novel and far more insightful probabilistic derivation and interpretation of results that were originally obtained using transform techniques from complex analysis. The successful interactions with France Telecom R&D on IP traffic measurement algorithms perfectly underscore the tremendous opportunities opened up by the unique combination of powerful mathematical techniques and the availability of large-scale real-life data sets. A delicate balancing act, which deserves and receives constant attention, arises from the fact that simple models, without nitty-gritty details, play a useful and entirely legitimate role in facilitating mathematical tractability, but should be used with caution in making inferences for more intricate, realistic situations.

Impact

The publication record of the team is excellent, both in terms of technical quality and productivity, especially in view of the relatively modest headcount. As mentioned above, the results have been published in flagship discipline-based journals such as Annals of Applied Probability and Random Structures and Algorithms, as well as prestigious application-oriented conferences, such as Mobicom and Performance. The extensive collaboration with France Telecom R&D serves as an effective and natural instrument for technology transfer.

Cooperations/coherence

The project is well coordinated and organized. Although the work on scaling laws for stochastic networks is somewhat separate from the effort on IP traffic measurement algorithms, both activities are sharply focused and well managed, and in fact exhibit some commonality in terms of the asymptotic methods used.

Strategic goals

As noted above, the proposed future research directions seem a natural progression of the ones for the previous four-year period, diminishing the work on bandwidth-sharing in IP networks and tree algorithms, and intensifying the effort on resource allocation mechanisms, with security issues as a particularly notable and mostly new avenue of research. The research agenda is well articulated and seems suitably aligned with the INRIA strategy. It strikes the right balance between theoretical challenges and practical applications, between ambition and risk, and between providing a sense of long-term strategic direction and leaving flexibility to take on emerging problems along the way.

Recommendations

The project is well positioned for continued success, on both fronts of fundamental research contributions and relevant application domains. In view of the relatively modest headcount, the team is potentially vulnerable to possible personnel changes, and some hiring might be help safeguard the viability.

I. TREC

Scope/vision

TREC is a joint INRIA-ENS project under the leadership of Professor F. Baccelli. It investigates networking fundamentals and their applications to modern network engineering (Internet, wireless). At the methodological level, it pursues fundamental research on communication networks control, modeling and performance analysis of wireless networks, stochastic network dynamics, and the development of mathematical tools based on stochastic geometry, random graphs, and spatial point processes. This research is always motivated by important problems in the corresponding areas of networks (both in the Internet and in wireless). It is very ambitious and targets at developing new mathematical tools for solving these problems. The results so far justify the risk taken by the group to tackle difficult problems using new mathematical tools. Among the examples we cite: stochastic geometry applications to wireless networks, the mean field analysis for wired and wireless networks, the development of stochastic network calculus for the analysis of rare events in networks. With respect to other INRIA projects in networking, TREC is centrally positioned by developing methodologies and network engineering solutions that are important and potentially of use to many of those other projects. TREC objectives as defined during the previous evaluation period have been by and large met.

TREC is an example of a truly outstanding project, doing world leading research according to the highest academic standards with high impact. Its team includes high-caliber researchers and is impressive in terms of the scope and quality of the results it has produced and the international visibility it has attained. The core team is rather small compared to other ComB team. It has remained more or less the same size over the last four years. The size is too small to have the impact it deserves. It is imperative that TREC be given any possible support to add a couple of high-quality researchers in its team and reach a critical mass of 5 full time researchers. TREC is clearly well aligned with INRIA's objective for quality research with high impact and is unique in pursuing successfully basic research with high but risky long-term impact.

We continue with detailed comments on specific topics.

Technical Excellence

TREC is characterized by technical excellence and by doing high-risk long term research. Probably it is uniquely positioned in this respect within INRIA. Professor Baccelli, who became a member of the French Academy of Science in 2005, is clearly a world-class leader in his field and has made seminal research contributions while working on this project on a continuing basis. An indication of his international recognition is the reception of the IBM Academic Award for 2004 and 2005.

The output of this team of researchers is also very impressive. Journal publications are in the top applied probability journals (the Annals of Applied Probability, the Journal and Advances of Applied Probability), and also in many journals of IEEE, JSIAM and ACM. Regarding conference publications, most of them are in the most prestigious conferences in the field (SIGCOMM, INFOCOM (9), SIGMETRICS (4), ITC (5)). Besides just purely academic achievement, the group has built key connections with industry leaders, like Alcatel and IBM, and has developed sophisticated software (N2N) for network dimensioning and design, which has led to the creation of a startup company (N2NSoft). Based on the above level of technical excellence and success, it is no surprise that TREC has attracted ample international visibility. We must also stress the impact of the work on IP networks. TREC got for two consecutive years (2004-5) the prestigious IBM Academic Award for its work on TCP overlays for multicasting. The resulting PhD thesis got the SPECIF prize and resulted in two patents!

There is also an important impact on academic teaching. There is an impressive number of new graduate courses taught by TREC members. There is a new book on the applications of stochastic geometry to wireless networks

now near completion.

Adequacy to INRIA's scientific strategy

TREC's achievements clearly align with the major goals of INRIA's scientific strategy to achieve scientific and technological breakthroughs at the top level worldwide in the design and master the future network infrastructures and communication services platforms by combining fundamental research activities with a strong involvement in technology transfer activities, industrial collaborations and company creations, by reinforcing partnerships with universities and engineering schools and by increasing the Institute's reputation at the international level.

Industrial relevance

TREC's research agenda is technically deep and insightful, as well as of high relevance and impact to network engineering. By exploring fundamental issues of network dynamics and control, it strengthens our understanding of such complex systems and it develops the much-needed scientific/systematic framework for new major technological innovations. TREC has been very eclectic in carefully choosing the research issues to address, striking a fine balance between theory and practice, and intellectual depth and engineering relevance. This well-coordinated combination has had big success on both fronts. A proof of the relevance of the work is the large number of grants and external funding of the group. There is impressive funding from the industry (FT, Thomson, Alcatel Lucent, IBM, Sprint, etc.), which is unusual for an academic group being so strong in theory like TREC, with less than 3 full time researchers on the average per year.

Impact (industrial-academic)

Many results of TREC are expected to have high impact. For instance, the max-plus algebraic studies of queuing/synchronization networks and the stochastic geometry applications on spatial traffic characterization for mobile communications have added significant tools in networking theory, with interesting engineering applications. Some excellent examples are the derivation of the spatial Erlang Formula and the self-organizing Gibbsian algorithms for WiFi LANs. We expect these results to eventually have high industrial impact and applicability. An existing impact proof already mentioned in the previous review is the work on TCP dynamics (Infocom '03 paper), which has led to a novel understanding of TCP attracting international attention, and resulted in the formation of the N2NSoft startup. We expect similar results to occur from the work on stochastic geometry regarding the topological design and control of wireless networks. We must also stress that the impact of the basic research done in TREC is expected to be rather long term. Another indication of the impact TREC has in the academic community on opening new research directions is the upcoming JSAC issue on Stochastic Geometry and Random Graphs for Wireless Networks where Prof. Baccelli is the guest editor. Stochastic geometry applied to networking research is becoming a mature field that has already attracted a large number of researchers internationally. This area of research was virtually unknown in the early 2000, prior to TREC publications.

Strategic goals

A first objective is to reach a critical mass of 5 permanent researchers. Assuming this is possible, the proposed objectives of TREC for the next four years are to do research in: i) overlay networks, ii) opportunistic and adaptive schemes in wireless (cellular, ad hoc) networks, iii) active network probing and the inverse problem, iv) go beyond the Poisson assumption in the stochastic geometry models, and v) start a new direction on message passing algorithms using random tree theory. All these directions are very important and challenging. They also naturally follow from the work performed so far in TREC. A sensitive issue is the ability to carry out research on all these fronts given the current small size of the team, and the risk of not achieving the critical mass of 5 researchers. Also one must take into account that it takes some time for a new member of the team to be productive given the sophistication of the mathematical tools used in TREC. One should be satisfied even if TREC achieves only a subset of the above proposed goals.

Cooperation/coherence

The TREC team is small and very well-integrated. Cooperation is high and coherent inside the team, still allowing each member appropriate independence to grow professionally and become a leader in his own specialty. Moreover, the project has built some good collaboration with leading researchers around the world as its publications record suggests. Regarding the researchers in the team, Dr. Blaszczyszyn has proved his research abilities by successfully continuing the work on stochastic geometry applications to wireless networks. Team member T. Bonald left for FT and A. Proutiere for Microsoft Research. D. Hong left for the startup N2Nsoft on 2003. All the above researchers demonstrated great potential and were rising stars in their fields at an early stage of their careers. It was unfortunate for TREC to loose them. To compensate that, Dr. Lelarge, a very promising young researcher, joined in 2006. C. Bordenave, a PhD student, is doing excellent and promising work collaborating with Dr. Lelarge and Prof. D. Aldous at UC Berkeley. A basic observation is that the size of TREC is rather small (compared to others in ComB) and has remained more or less at the same size since its previous evaluation cycle. It clearly suffers from not operating above the critical mass of full time researchers. It is imperative that TREC be given by INRIA any possible support to add at least two high-quality researchers to its team and reach the critical mass of 5 full time researchers.

Recommendations

In conclusion, we believe that the TREC project is a clear success story in INRIA's research program generating international visibility via top level quality research with practical impact. The project has gathered impressive momentum and opened up new methodological and technological research avenues to pursue. Even more impressive is that fact that this success has been achieved by a small team of high-caliber researchers, some of whom have since left the team to join the industry research labs. The important issues that we see are the following.

(a) First, priority must be assigned to growth of the group to reach a critical size. Although the successful transfer of former members to excellent positions in industrial laboratories is a clear proof of TREC quality and reputation of its research, it has proved to have a negative effect on the size of the group. One solution might be to offer permanent research positions to current brilliant PhD students once they graduate, until TREC reaches a comfortable size. Since it is well understood that researchers with the expertise required to contribute in TREC are hard to find, it makes it imperative for the INRIA management to assist the project in this direction. We must reiterate that TREC should reach a critical mass of at least 5 full time researchers. We must note that this may not be an easy task. It might prove to be harder than in other projects due to the high theoretical skills required. For such highly skilled researchers industry may be able to offer more attractive and competitive position than INRIA. Hence the INRIA management must find some ways to compensate for that.

(b) The second issue is that TREC should keep its efforts in disseminating the research done within the group. It is hard for such a small group of researchers to be very successful in all fronts, namely: exploring new research areas, attracting funding, teaching courses on new subjects, disseminating the research results among peers. We focus on dissemination because the new tools for modeling and analyzing complex systems developed within TREC look very promising for adoption by the wider research community. This is perhaps a unique strength of TREC, namely, to be a constant source of innovative research by aggressively selecting new advanced mathematical concepts and formalisms and adapting them successfully to networking research and allowing hard problems to become tractable. This is made possible by the unique mathematical talents of its researchers and its team leader. Hence helping the rest of the research community to assimilate these results and methods is of extreme importance. TREC has done a first rate job in this direction so far. The team might achieve even better dissemination by organizing short courses or tutorials on these topics for other researchers in the future.

(c) It is clear that up to now Prof. Baccelli has played a central role in all the research activities in TREC. We hope that in the future as TREC grows to critical mass, more researchers will be able to play a central role and assist Prof. Baccelli in pursuing and managing these new research directions. This will be a sign of maturity for TREC and an indication of its long-term sustainability. A positive sign in this direction is that Dr. Blaszczyszyn and Dr. Lelarge seem very promising and capable in their areas of research, and are expected to have significant impact in the evolution of TREC by growing their individual areas of research.

IV. General Conclusions

The evaluation Panel assessed the nine projects of ComB (ARES, DISTRIBCOM, HIPERCOM, MADYNES, MAESTRO, MASCOTTE, PLANETE, TREC, RAP). The general conclusions that emerge from the project evaluations is that the ComB program as a whole is on track. Performance over the past four years has been outstanding. The goals set forth by the project leaders at the beginning of the four year cycle have been by and large accomplished; the objectives formulated for the next four years are realistic and consistent with INRIA strategic objectives. Reviewing the general criteria established by INRIA for this external evaluation, the Panel has determined that the Program has successfully met these criteria:

Scientific excellence

Assessment is based on the scientific originality of the project-team, the results obtained and anticipated, the difficulties that are overcome or that may be predicted, successful and hoped-for innovations and the methods used. All these elements are assessed against the goals previously set out. The new scientific objectives of the project-team for the next four years should be presented, accompanied by a timescale; the priorities and any necessary reorientations should be described.

All the teams have produced ample evidence of originality and excellence in their research objectives, through the volume and quality of their publications as well as through the transfers of their results (software, tools, theory, methodologies) to the International Community.

Adequacy to INRIA's scientific strategy

The purpose here is to evaluate the project-team contributions to the accomplishment of INRIA strategies, together with its contacts and collaborations at national and international levels. Particular attention is drawn to the domains that are not covered and possible ways of covering them with existing project-teams.

The teams have all achieved in various degrees international visibility in their respective fields by carefully and effectively pursuing directions of competitive excellence. Moreover, many of the projects have proactively established collaborations with centers of excellence throughout the world. Through the ComB program, INRIA is maintaining a very high profile in the international arena in the area of Networking and Telecommunications.

Industrial transfer and partnerships

The applications, particularly the industrial applications, of the project-team are examined, together with the partnerships to which it has contributed. The contractual activities are qualitatively and quantitatively assessed, bearing in mind that apart from the contract itself it is important – though not always easy – to be able to see the impact of the project-team's research work.

All the teams reviewed in this phase have stressed in their reports and presentation their interest in the transfer of their research results to industry. The success in this transfer has varied from project to project, as expected, since INRIA's strategy is to maintain a presence in all areas of Networking and Telecommunications, from fundamental to applied research. Thus, some projects are more amenable to commercialization than others. This panel was pleased to notice that a substantial part of the research budget is contributed by Industry and by EEC projects, underlying the importance of transfer and international collaboration.

Manpower and means

The development of the reputation in the scientific community of the researchers taking part in the project-team is examined. Changes in project-team personnel are considered (number of researchers appointed, proportion of senior researchers and directors, number of PhD students, etc.) and, in particular, mobility of tenured researchers towards the universities and industry.

In INRIA as everywhere else the success of the research project is entirely determined by the leadership of the research director, by the excellence of the individual researchers and by the ability of them to work as a team. During the interviews, this Panel was able to observe first hand the leadership and team spirit that make these projects tick and deliver. Taken individually, the researchers all have a well established track record and reputation. They make an effort to improve INRIA visibility not only through publications but also through active participation in Conference Committees, Journal Editorial Boards and Teaching Engagements in top academic Institutions. Most

of the teams are very ”lean”, with relatively few permanent members. In fact, most teams complement their research force with part timers from academia and with interns from local PhD programs. To this end, we recognize the important role of the INRIA projects in educating young researchers, witnessed by the numbers of PhD’s produced by the projects over the past four years and the numbers of PhD students currently employed.

In closing, the panel notes that ComB - Networks and Telecommunications - continues to play an important role within INRIA in the development of novel ideas and concepts for current and future information and communication technologies. In particular, the panel acknowledges and praises the entrepreneurship and proactive effort of the individual projects in setting their goals so as to exploit their competitive edge, maximize their research outputs and strike the best balance in collaborating with industry, academia and other INRIA projects. The relative autonomy and self management of the teams during each four year period is an important part of the INRIA culture and has worked extremely well so far. The panel strongly recommends that this culture be preserved. At the same time it recommends that, for the further positioning of INRIA as a world leader in the field, the overall strategic goals of the collection of these projects be made more explicit and that the research on common technology issues be actively cross-referenced across project boundaries.

A. APPENDIX A: Evaluation Criteria

List of evaluation criteria used by the panel as a guideline

- Vision. Overall expression of long-term goals.
- Scope. Significance (overall), contribution to INRIA strategy, overall impression (of past performance and role in wider scientific context), objectives (current and future), what are the peer groups.
- Technical excellence. Assessment of technical quality of the work, originality, depth, innovations in methodology or results. Scientific results (publications), patents, PhD’s. Strengths and weaknesses.
- Importance/relevance. Most striking results, success stories, interplay between theory and practice, innovations used by others outside project and outside INRIA, connection to state-of-the-art.
- Impact. Publication/conference record (qualitatively), usefulness for science-industry economy- society, spin-offs, commercial products, books, software produced, technology transfer, international visibility.
- Cooperations/coherence. How is the project functioning as a project, organization and management, cooperation with other project-teams, activity in internal and external contacts, is effort commensurate with importance.
- Strategic goals. Directions for future research, contribution to INRIA strategic plan, expected results, risk analysis, innovations, links/plans with new applications and partnerships, relevance to INRIA.
- Recommendations. Analysis of major problems (if any), issues that need attention, scientific advice, opportunities and threats.