

FET Proactive initiative

EVERGROW
Integrated project



<http://www.evergrow.org/>

001935

Ever-growing global, scale-free networks, their provisioning, repair and unique functions

Review Report N°: 3

Covering period: 01/01/2006-31/12/2006

Report Preparation Date: 06/03/2007

Report: Restricted

Contract Start Date: 01.01.2004 Duration: 48 months

Project Co-ordinator: SICS (Seif Haridi) & Hebrew University (Scott Kirkpatrick)



Place of the Review: ETH, Zurich, Swizerland

Date of review: March 5th & 6th 2007

Reviewers:

Prof. Mark Crovella
Boston University, MA, USA

Dr, Hab Anne-Marie Kermarrec
INRIA, Rennes, France

Prof. David Dean
Université Paul Sabatier, Toulouse,
France

Prof David Lane
University of Modena
Italy

European Commission

Project officer: Fabrizio Sestini

OVERALL ASSESSMENT OF THE PROJECT

Evergrow is a large and ambitious project focusing on the future of the Internet, which is viewed in this project as a complex system. EVERGROW is part of a cluster of FET projects of the Complex Systems initiative. The reporting period addressed by the review was the third year of the four-year project.

The vision of the project, updated over the 2nd year of the project is to provide the ultimate Google/Akamai/Raid/Gnutella. Evergrow works towards these objectives along several lines

- measuring and modelling the current Internet infrastructure and traffic
- designing relevant peer to peer overlay networks and algorithms for the considered applications mapping into the vision
- developing new tools to study complex systems, inspired partially by physics based approaches, which are then exploited for measurement analysis and modelling purposes.

Over the past year, the project has been consolidated, and the review team was impressed by the high quality of the research results achieved. The team was also impressed by the implementation effort in measuring the Internet and providing relevant visualization tools. If we were to highlight a few achievements:

- the results achieved with respect to DIMES and ETOMIC individually and also in terms of their integration. This is the largest measurement effort of its type so far conducted and the measurements obtained to date are unique in their precision and novelty.
- The work in the area of peer to peer overlay networks is highly innovative (Best paper award at IPDPS 2006, Julia protocol, Oscar system)
- The research on survey propagation is world-leading. We acknowledge the innovation of the work on finite length stochastic codes and other aspects of this WP that, in many cases, have been achieved well ahead of schedule
- The plan for the work in integrating measurements with modelling and emulation environments goes beyond expectations.

We also acknowledge the specific effort to integrate SPs over the past year. Many workshops involving partners from different SPs have been organized and have led to some successful collaboration and to applications of results of SPs in other SPs. For example SP3 exploited the results obtained in SP4, while SP4 has also be driven and guided by the requirements and better understanding of SP3. This effort started during the 2nd year and has been successfully carried on over the review period. This synergy remains crucial for the last year of the project where we expect further integration.

Overall, the review outcome is extremely positive. The reviewers recognized the scientific achievements and the innovation of the approach and recommend the continuation of the project together with a few adjustments.

1 PROJECT ACHIEVEMENTS and FUTURE PLANS

1.1 Work carried out in the previous reporting period:

Scientific and Technical work

On the scientific side, the review team judged most parts of the project to be on schedule: SP1 is on or ahead of schedule. SP2 is now on schedule, and some aspects of SP4 are somewhat ahead of schedule. Most sub-projects have achieved high-quality scientific results and the results of the synergy between sub programs have proven effective.

SP 1: Measurement and Modelling

The measurement and modelling subproject has made considerable gains in the past year.

The DIMES agent and system have proved to be reasonably stable and have accelerated their rate of penetration and data collection. Further, the project has worked hard to make the system and its data widely available to researchers. The project has implemented a new experiment planner that allows outside researchers to run experiments on DIMES, and it has implemented a web-based interface to the DIMES data.

The ETOMIC effort has also solidified. The system is running in a stable mode. An available bandwidth measurement capability has been added as well as large-scale inference of network-internal delay distributions.

Further, there is good progress on unifying the ETOMIC and DIMES systems allowing their techniques to be used together.

Reviewers were impressed with the accomplishments during the reporting period, but also had a number of recommendations. It is recommended that ETOMIC not cast itself as “traffic measurement” since traffic is not being measured directly. Furthermore, the extraction of key parameters from DIMES and ETOMIC for construction of the Internet model is considered an unsolved and challenging task.

The precise way that ModelNet / GODS will be used as a testbed for P2P designs is still unclear. This integration between WP1c and WP3 still needs development. Furthermore, the use of measurements from ETOMIC for parameterizing such a model needs further consideration. Use of aggregate delay statistics across all links, and use of inferred network-internal properties as inputs to ModelNet will need to be carefully examined.

Finally, the model of Internet evolution is an ambitious goal and the reviewers urge the team to address this goal energetically.

1.1.1 SP2:

The infrastructure for experimentation and analysis is now substantially in place and functioning satisfactorily. After encountering substantial problems in the first year, the set of 8 clusters located at consortium institutions has now been successfully integrated with Everlab, a private PlanetLab, supplemented by high performance computing tools, which supports both network experimentation and computation. With the addition of the opensource centralized monitoring Ganglia software and their own custom built Everstats resource reporting system, administrators can track system usage and identify network and hardware problems. The whole system constitutes a significant achievement for Evergrow, and the project should over the next year develop a plan for how it might open it for use to outside researchers, provide documentation that will make it easier for the next such system to be constructed, and most importantly provide for a continuing existence for the infrastructure and the Virtual Network Observatory that will be based upon it. This last issue is particularly delicate and difficult, since warranties will expire by the end of 2007 on the clusters, and blade disks have proven to be very unstable.

The Virtual Network Observatory is potentially the most valuable legacy of Evergrow, since it could provide to researchers throughout the world valuable information on the current structure and functioning of the physical Internet and how these are changing in time. It will be interesting to see how well the VNO serves the needs of the Evergrow group that intends to build a model on Internet evolution in the coming year, and how successfully the data from the VNO can be integrated into the ModelNet/GODS testbed environment that Evergrow is putting in place, to provide realistic and useful parameters for the underlying network in the overlay network emulations that ModelNet/GODS are to support.

The three deliverables from SP2 are acceptable. It is hoped that the data analytic and modelling capabilities of the CLAES group be integrated into the challenging Internet evolution modelling task.

SP3: Peer to peer overlays and service architecture

The SP3 work targets the P2P overlay networks and some specific applications. The spectrum is still quite large (search, content distribution, storage...). This past year, the main results are related to

- The Julia content distribution protocol, started in 2005 (Ultimate Akamai)
- The myriad back-up systems (Ultimate Raid)
- The Oscar system
- The DHTtrie system

As mentioned in previous reviews, the outcome of this SP is a collection of implementations, of algorithms developed mostly (except for Julia) on structured P2P overlays (Chord or DKS). This aspect seems to confirm the fact that no single overlay is suitable for all problems and the consortium is aware of this and is working on a way to integrate this point into further work.

The reviewers felt that an integration, not implementation-wise, but concept-wise should be addressed in the future. An orthogonal concept seems to be the use of message passing and belief propagation algorithms.

As far as the deliverables are concerned, some of them rely again heavily on papers, we would advise the consortium to work on more integrated deliverables (backed up by papers obviously) for the final review. However, the review team acknowledges the quality of both the oral presentation, which were of high quality, as well as the research carried out.

SP4: Theory of coding and algorithms

Message passing and belief propagation techniques and finite length coding have been studied from a theoretical point of view. This research is world- leading and has led to a significant production of high quality research papers, conference presentations and a patent. The algorithms inspired by SP4 research have also started to be adopted by other groups. Message passing has been implemented by SP3 to analyse internet measurement data and has substantially reduced the computational time associated with the analysis. The impression that the reviewers got at the review meeting was that the other SP s are increasingly convinced by the usefulness of SP4 developments and that the techniques developed by SP4 should see more extensive application in the final stage of the project.

The SP4b work package on finite-length coding has produced its final deliverable this year, that is to say a year early. This is due in part to the departure of the leader of this WP for a position in the US. The remaining resources will be allocated to SP4a for the last year of the project. The work of SP4b contains a significant number of contributions. An important point is that a number of results obtained via statistical physics methods have been made rigorous (using quite recent results in mathematical physics). Despite the fact that this added rigour does not necessarily lead to new results, it is important in the sense that it will play a role in convincing the computer science community of both the validity and usefulness of physics-inspired approaches. Significant progress has also been made in the application of belief propagation and new inference techniques to coding for peer to peer contents distribution; these new methods present a considerable improvement with respect to traditional methods. The theory of finite length codes has also been studied and a scaling function, similar to those arising in the theory of phase transitions, describing the block error coding rate about the threshold for reliable communication has been studied.

Resources

As far as resources are concerned, the main issue relates to the fact that although the computer clusters are up and running, they are unfortunately underused. The review team suggests that some provision for opening it to researchers outside the project, under the appropriate control, should be considered.

Status of deliverables

The deliverables have been produced on time. The quality of the deliverables is good overall, although some of them could benefit from typo and spell checking. The review team found the deliverables satisfactory, and we suggest they be accepted without any modification.

1.2 Work planned for the next 18-month period:

The plan for the next 18 months proposed by the consortium during the review meeting was considered satisfactory.

Scientific and technical aspects:

Comments relevant to specific subgroups

- SP1: Greater clarity is needed with respect to the ModelNet / GODS approach to bringing results from SP1 into SP3. Further, future plans for continued operation of the cluster, including continued access to DIMES and ETOMIC data and systems need to be addressed in the coming year (before the end of the project).
- SP2: Plans for continued operation of the cluster, archiving strategy, access by the research community to data and experimental opportunities in the VNO all need to be developed and publicized. In addition, documentation that can help other groups benefit from what Evergrow learned in constructing the Everlab environment would be useful. The formation of a team with requisite data analytic and modelling capabilities, including CLAES, should be completed soon, in order to arrive at a successful outcome for the Internet evolution model.
- SP3: the issue of integrating the various systems from a conceptual point of view has been discussed. The main objective that the review team is looking forward to seeing implemented during the last year is related to the performance evaluation of P2P exploiting Internet measurements. This translates into using ModelNet over the realistic Internet topologies coming from the DIMES/TOMIC measurements and running and evaluating the P2P systems of SP3 in that context.
- SP4: During the final stages of the project, given that this SP is to some extent ahead of schedule in terms of theoretical development, it is hoped to see further impact of this group's research on the activities of the other SPs and towards the global objectives of the program

Non-scientific activities:

Given the interaction that took place between subprograms over the past year, the reviewers are confident that the coming year will be as successful.

2 CONSORTIUM PARTNERSHIP

The degree of interaction between partners really initiated last year increased again in 2006 and this synergy is reflected in the work produced by the consortium. The use of results from

each SP in the others SPs is now clear and efficient. The review team once more acknowledges the individual excellence of each of the SPs as is confirmed by their first class academic records.

3 PROJECT MANAGEMENT AND CO-ORDINATION

The review team was satisfied by the way the project has been managed and coordinated over the reviewing period. The main problem concerns the hardware chosen for the cluster, which led to highly unstable environment prone to failures. The consortium has been misled by their suppliers in respect: they were assured that the hardware was enterprise-level equipment; however it has turned out that it is well below the level expected.

4 USE AND DISSEMINATION OF KNOWLEDGE

The review team judged the efforts made for the use and dissemination of knowledge to be satisfactory. However, the review team thinks that the quality of DIMES (and more specifically the large-scale Internet mapping) as well as the visualisation and analysis packages of DIMES and ETOMIC could be advertised even further. In addition, some attempts to describe and disseminate material that will enable other groups to construct Everlab-like experimentation and computation infrastructures should be carried out.

The review team considers that the impact of the project might be important to understand the future Internet infrastructure which heavily depends on the current one as well as on the potential applications.

5 CONCLUSION and SUMMARY of RECOMMENDATIONS

Good to excellent project (The project has fully achieved its objectives and technical goals for the period and has even exceeded expectations)

Recommendation

- the project should continue with the following modifications (technical or administrative)
- The consortium should address the under-use of the cluster by encouraging outside use
 - The consortium should formulate plans for continuation/strengthening of the virtual network observatory beyond the end of the project.

Name(s) of the reviewer(s):

Mark Crovella

David Dean

Anne-Marie Kermarrec

David Lane

Date: March 26th, 2007

Signature(s):

6 APPENDICES

1. Status of project reports and deliverables (submitted/delayed, accepted/rejected/to be modified)
2. List of participants
3. Agenda

Appendix 1

Status of project reports and deliverables.

Deliverable number	Title	Status (submitted/delayed)	Accepted/Rejected/To be modified	Deadline for re-submissions	Remarks
D.1a.3	Web-based interfaces for query access to live DIMES databases	15/02/07	Accepted		
D.1b.3	Spatio-temporal periodic mapping of the European Internet traffic	15/02/07	Accepted		
D.1c.1	Critical Network Properties for Peer-to-Peer Systems Performance Evaluation	19/02/07	Accepted		
D.2a.2	Virtual Network Observatory--first report	15/02/07	Accepted		
D.2c.1	Documentation of data content	15/02/07	Accepted		
D.2d.3	Working prototype system to label and access temporal data from DIMES	15/02/07	Accepted		
D.3k.2	Report on extend maintenance, routing, and task-specific algorithms for overlay networks	15/02/07	Accepted		
D.3l.2	Second report on higher-level services on top of overlay networks	15/02/07	Accepted		
D.3m.2	Final report on distributed storage and content-distribution	15/02/07	Accepted		
D.3n.2	A Detailed Analytical Study of the Effects of Churn and Maintenance on Overlays	15/02/07	Accepted		
D.4b.4	Finite-length coding and application to multi-user systems	15/02/07	Accepted		

Appendix 2

List of participants at the review.

- Yuval Shavitt, (SP1,WP1a) Tel Aviv University
- Gabor Vattay, (SP1.WP1b) Eotvos University and Collegium Budapest
- Vasilios Dargilianis (SP1, WP1c and SP3) EPFL
- Elliot Jaffe (SP1, WP2a), HUJI
- Scott Kirkpatrick (WP2c), HUJI
- Mohammed El Beltagy (WP2d), CLAES (Cairo)
- Danny Dolev (SP3) HUJI
- Danny Bickson (SP3) HUJI
- Supriya Krishnamurthy (SP3) KTH and SICS
- Karl-Filip Faxen our administrator (SICS)
- Federico Ricci-Terzenghi (SP4) Rome
- Lenka Zdeborova (SP4) Paris
- Valerie van Kerrebroeck (SP4) Rome
- Karl Aberer (EPFL) (SP3)

Appendix 3

Agenda of the meeting

Monday 10 am - 1 pm

- DIMES and ETOMIC measurement activities 2 short talks + demo 1 hr
- WP1c integrating measurement with P2P 1/2 hr status report
- SP2 Virtual Observatory 2 talks on computing platform and on tools for analysis + domes 1 hr.
- plus time for discussion and additional demo

Monday 2:30 - 4:30 pm

- SP3 P2P architectures and applications 2 talks
- SP4 message passing for scalable algorithms 2 talks

break

Monday 5 - 6 pm

- Administrative and financial details
- Overview of 2007 final year plan

Tuesday 9-10

- Details of 2007 final year plan
- 10-11 Reviewers' caucus
- 11-12 Feedback session