Rapport de fin de Projet
No de Projet: ANR-05-SSIA-0016-01
PROJET ANR SOGEA.

30 Septembre 209

A. Identification

Programme-Année	SSIA SOGEA 2005
Projet (acronyme)	SOGEA / ARA-SSIA-2005
Coordinateur du projet (société/organisme)	Olivier Bournez (INRIA, LORIA, Ecole Polytechnique)
Période du projet (date début - date fin) contractuelle)	1ier Décembre 2005 au 1ier Décembre 2008 étendue au 30 juin 2009

Rédacteur de ce rapport	
civilité, prénom, nom	Mr Olivier Bournez
téléphone	01 69 33 40 78
adresse électronique	Olivier.Bournez@polytechnique.fr

B. Pour les projets partenariaux, rappel des livrables ou jalons alloués aux partenaires pour l'ensemble du projet

Le document scientifique soumis ne promettait pas formellement de délivrables, mais faisait une liste d'un certain nombre de tâches à effectuer.

Nous reprenons ici les informations de la partie scientifique du projet: Intitulés (partenaires entre parenthèses).

- Tache 1.1: " Notions de stabilité " (LORIA, LRI, PRISM)
- Tache 1.2: " Modèle du routage inter-domaine " (LORIA, PRISM)

- Tache 1.3: " Competitive Self-Stabilization " (LRI, LORIA)
- Tache 1.4: " Aspects Dynamiques " (LORIA, PRISM)
- Tache 1.5: " Techniques de preuves associées " (LORIA, PRISM)
- Tache 1.6: " Autostabilisation résistante aux comportements maléfiques " (LRI)
- Tache 2.1 : " Algorithmes distribués pour le routage interdomaine avec aspects économiques". (LRI, LORIA, PRISM)
- Tache 2.2 : " Algorithmes séquentiels pour réseaux de capteurs " (LORIA, LRI)
- Tache 2.3 : " Algorithmes distribués avec garanties dynamiques pour le routage interdomaine (LRI, PRISM).
- Tache 2.4 : " Algorithmes distribués malicieux " (LRI, PRISM)
- Tache 2.5 : " Algorithmes distribués d'autostabilisation compétitive " (LORIA, LRI, PRISM)
- Tache 3.1 : " Résultats de complexité " (LORIA, LRI)
- Tache 3.2 : " Méthodes d'approximation " (LORIA, LRI, PRISM)
- Tache 3.3 : " Vérification efficace d'équilibres " (LRI)
- Tache 3.4 : " Simulations " (LORIA, PRISM)

Autrement dit:

Nancy	1.1	1.2	1.3	1.4	1.5		2.1	2.2	2.3		2.5	3.1	3.2		3.4
Orsay	1.1		1.3			1.6	2.1	2.2	2.3	2.4	2.5	3.1	3.2	3.3	
Versailles	1.1	1.2		1.4			2.1		2.3	2.4	2.5		3.2		3.4

Calendrier:

0-6 mois	Tâches 1.1, 1.2, 1.3, 1.4, 1.5, 1.6
6-12 mois	Tâches 1.1, 1.2, 1.3, 1.4, 1.5, 1.6
12-18 mois	Tâches 1.1, 1.2, 1.6, 2.1, 2.2, 2.4, 2.5,
18-24 mois	Tâches 2.1, 2.2, 2.3, 2.4, 2.5, 3.1, 3.2, 3.3, 3.4
24-30 mois	Tâches 2.1, 2.2, 2.3, 2.4, 2.5, 3.1, 3.2, 3.3, 3.4
30-36 mois	Tâches 3.1, 3.2, 3.3, 3.4

Coordinateurs: Olivier Bournez pour le LORIA, Dominique Barth pour le PRISM, Sylvie Delaët pour le LRI.

C. Rapport Factuel

Ce rapport rassemble des éléments nécessaires au bilan du projet et plus globalement permettant d'apprécier l'impact du programme à différents niveaux. Il est constitué de deux tableaux et de deux listes rassemblant des indicateurs quantitatifs, bibliométriques et de personnel.

C.1. Le tableau de résultats

décompte divers éléments de bilan comme indiqué ci-dessous

Nombre de Publications

- International:
  - Chapitres de livres:
    - Monopartenaire: [10, 42]
    - Multipartenaires: ∅.
  - Habilitations à Diriger les recherches:
    - [9, 41]
  - Articles acceptés dans des revues à comité de lecture:
    - Monopartenaire: [19, 38, 8, 37]
    - Multipartenaires: [17].
  - Communications internationales:
    - Monopartenaire: [34, 43, 27, 39, 25, 23, 15, 36, 35, 31, 30, 28, 18, 20, 21, 22, 32, 29, 27]
    - Multipartenaires: [2, 11, 5]
- France: ∅.
  - Communications France:
    - Monopartenaire: ∅.
    - Multipartenaires: [1]
- Actions de diffusion: ∅.
- Autres:
  - Monopartenaire: rapport de stage [33], Rapport de recherche [3, 4, 24, 6, 7, 26, 40]
  - Multipartenaires: [12, 14, 13]

Autres retombées (voir en particulier celles annoncées dans l'annexe technique) : *Ce tableau dénombre et liste les brevets nationaux et internationaux, licences, et autres éléments de propriété intellectuelle consécutifs au projet, du savoir faire, des retombées diverses en précisant les partenariats éventuels.*

Nature	Commentaire
Brevets nationaux	Néant
Brevets internationaux	Néant
Autres	Néant

C.2. Le tableau de personnels fait le point sur les emplois mis en jeu et consécutifs au projet.

- Doctorants: $\emptyset$.
- Post-Docs: 1×12 mois (Mariusz Rokiki, démissionnaire pour raisons familiales, puis Stéphane Devismes) + 1×12 mois (Marc-Antoine Weisser)
- Ingénieur en CDD: 1×24 mois (Octave Boussaton)
- Stagiaires: Différents stagiaires de très courte durée.

Devenir des personnes employées en CDD sur le projet

- Octave Boussaton: actuellement inscrit en thèse.
- Mariusz Rokiki, actuellement post-doctorant.
- Stéphane Devismes: actuellement maître de conférence à VERIMAG, Université de Grenoble.
- Marc-Antoine Weisser: actuellement maître de conférences à Supélec, Gif sur Yvettes.

C.3 La liste des publications et communications

répertorie les productions résultant des travaux effectués dans le cadre du projet. On suivra les catégories du tableau C.1 en suivant les normes éditoriales habituelles de présentation des publications et brevets. En ce qui concerne les conférences, on spécifiera les conférences invitées.

L'ensemble des publications est disponible sur
`http://sogea.loria.fr/WIKI/pmwiki/pmwiki.php?n=Main.RelatedPublications`

La classification suivant le tableau C.1 se trouve dans la section C.1.

References

[1] D. Barth, O. Bournez, O. Boussaton, and J. Cohen. Convergences et dynamiques du routage dans les réseaux. In *Journées Pôle ResCom*, September 2007.

[2] Dominique Barth, Olivier Bournez, Octave Boussaton, and Johanne Cohen. Distributed learning of wardrop equilibria. In *Unconventional Computation 2008, UC 2008*, volume 5204 of *Lecture Notes in Computer Science*, pages 19–32, Vienna, Austria, August 25-28 2008. Springer.

[3] Dominique Barth, Olivier Bournez, Octave Boussaton, and Johanne Cohen. A dynamical approach for load balancing. Technical report, LORIA/INRIA, 2008. Submitted. Available on `http://www.lix.polytechnique.fr/~bournez/load/Soumis-Octave-Fev-2008.pdf`.

[4] Dominique Barth, Octave Boussaton, Olivier Bournez, and Johanne Cohen. Selfish distributed routing and convergence to nash equilibria in wardrop networks. Technical report, LORIA-INRIA, 2008.

[5] Dominique Barth, Johanne Cohen, Loubna Echabbi, and Chahinez Hamlaoui. Transit price negotiation: a combined game theoritic and distributed algorithmic approach. In *International Conference on Network Control and Optimization (EuroFGI NET-COOP'2007)*, volume 4465 of *Lecture Notes in Computer Science*, 2007.

[6] Dominique Barth and Chahinez Hamlaoui. Impact of the selfishness of pricing strategies on the qos provisioning in interdomain networks. Technical report, PRISM, 2007. Submitted.

[7] Dominique Barth and Chahinez Hamlaoui. Modèle degestion de la qos dans les réseaux interdomaine basé sur les jeux répété. Technical report, PRISM, 2007. Submitted.

[8] Joffroy Beauquier, Sylvie Delaët, Shlomi Dolev, and Sébastien Tixeuil. Transient fault detectors. *Distributed Computing*, 20(1):39–51, 2007.

[9] Olivier Bournez. *Modèles Continus. Calculs. Algorithmique Distribuée.* Hdr, Institut National Polytechnique de Lorraine, 7 Décembre 2006.

[10] Olivier Bournez and Manuel L. Campagnolo. *New Computational Paradigms. Changing Conceptions of What is Computable*, chapter A Survey on Continuous Time Computations, pages 383–423. Springer-Verlag, New York, 2008.

[11] Olivier Bournez, Philippe Chassaing, Johanne Cohen, Lucas Gerin, and Xavier Koegler. On the convergence of a population protocol when population goes to infinity. In *Physics and Computations, Worshop of Unconventional Computation 2008, UC 2008*, Vienna, Austria, August 25-28 2008.

[12] Olivier Bournez and Johanne Cohen. Learning equilibria in games by stochastic distributed algorithms. Technical Report abs/0907.1916, CoRR, 2009.

[13] Olivier Bournez and Johanne Cohen. Learning equilibria in games by stochastic distributed algorithms. *CoRR*, abs/0907.1916, 2009.

[14] Olivier Bournez and Johanne Cohen. Stochastic learning of equilibria in games: The ordinary differential equation method. Technical report, PRiM, Université Versailles, 2009.

[15] Olivier Bournez and Emmanuel Hainry. On the Computational Capabilities of Several Models. In *Machines, Computations and Universality (MCU'2007)*, volume 4664 of *Lecture Notes in Computer Science*. Springer, September 10-13 2007.

[16] Johanne Cohen, Anurag Dasgupta, Sukumar Ghosh, and Sébastien Tixeuil. An exercice in selfish stabilization. Technical report, LORIA/LRI, 2008. Submitted.

[17] Johanne Cohen, Anurag Dasgupta, Sukumar Ghosh, and Sébastien Tixeuil. An exercise in selfish stabilization. *ACM Transactions of Adaptive Autonomous Systems (TAAS)*, 2008.

[18] Praveen Danturi, Mikhail Nesterenko, and Sébastien Tixeuil. Self-stabilizing philosophers with generic conflicts. In Ajoy K. Datta and Maria Gradinariu, editors, *Eighth International Symposium on Stabilization, Safety, and Security on Distributed Systems (SSS 2006)*, Lecture Notes in Computer Science, pages 214–230, Dallas, Texas, November 2006. Springer Verlag.

[19] Praveen Danturi, Mikhail Nesterenko, and Sébastien Tixeuil. Self-stabilizing philosophers with generic conflicts. *ACM Transactions of Adaptive and Autonomous Systems (TAAS)*, 2008.

[20] Anurag Dasgupta, Sukumar Ghosh, and Sébastien Tixeuil. Selfish stabilization. In Ajoy K. Datta and Maria Gradinariu, editors, *Eighth International Symposium on Stabilization, Safety, and Security on Distributed Systems (SSS 2006)*, Lecture Notes in Computer Science, pages 231–243, Dallas, Texas, November 2006. Springer Verlag.

[21] D.Barth, L.Echabbi, and S.Hamlaoui. Transit price negotiation: Decentralized learning of optimal strategies with incomplete information. In *Next Generation Internet Networks NGI'2008*, 28-30 April 2008 2008. Available on Digital Library IEEE Xplore.

[22] Michel de Rougemont and Adrien Vielleribière. Approximate data exchange. In *11th International Conference on Database Theory (ICDT'2007)*, volume 4353 of *Lecture Notes in Computer Science*, pages 44–58, Barcelona, Spain, January 2007 2007.

[23] Sylvie Delaët, Stéphane Devismes, Mikhail Nesterenko, and Sébastien Tixeuil. Brief announcement: Snap-stabilization in message-passing systems. In *Principles of Distributed Computing (PODC 2008)*, August 2008.

[24] Sylvie Delaët, Stéphane Devismes, Mikhail Nesterenko, and Sébastien Tixeuil. Snap-stabilization in message-passing systems. Research Report 6446, INRIA, 02 2008.

[25] Sylvie Delaët, Stéphane Devismes, Mikhail Nesterenko, and Sébastien Tixeuil. Snap-stabilization in message-passing systems. In *International Conference on Distributed Systems and Networks (ICDCN 2009)*, January 2009.

[26] Sylvie Delaët, Partha Sarathi Mandal, Mariusz Rokicki, and Sébastien Tixeuil. Deterministic secure positioning in wireless sensor networks. Research Report 6326, INRIA, October 2007.

[27] Sylvie Delaët, Partha Sarathi Mandal, Mariusz Rokicki, and Sébastien Tixeuil. Deterministic secure positioning in wireless sensor networks. In *Proceedings of the ACM/IEEE International Conference on Distributed Computing in Sensor Networks (DCOSS 2008)*, Lecture Notes in Computer Science. Springer-Verlag, June 2008.

[28] Stéphane Devismes, Sébastien Tixeuil, and Masafumi Yamashita. Weak vs. self vs. probabilistic stabilization. In *Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS 2008)*, Beijin, China, June 2008.

[29] E. Fischer, F. Magniez, and M. de Rougemont. Approximate Satisfiability and Equivalence. In *Proceedings of 21st IEEE Symposium on Logic in Computer Science (LICS'2006)*, pages 421–430, 2006.

[30] Maria Gradinariu and Sébastien Tixeuil. Conflict managers for self-stabilization without fairness assumption. In *Proceedings of the International Conference on Distributed Computing Systems (ICDCS 2007)*, page 46. IEEE, June 2007.

[31] Fabíola Greve and Sébastien Tixeuil. Knowledge connectivity vs. synchrony requirements for fault-tolerant agreement in unknown networks. In *Proceedings of IEEE International Conference on Dependable Systems and networks (DSN 2007)*, pages 82–91. IEEE, June 2007.

[32] Sébastien Hémon, Michel de Rougemont, and Miklos Santha. Approximate nash equilibria for multi-players games. In *First International Workshop on Algorithmic Game Theory (SAGT'08)*, Paderborn, Germany, April 30-May 2 2008 2008.

[33] Xavier Koegler. Protocoles de population continus. Stage de prédoctorat ENS Paris, 2006.

[34] Fredrik Manne, Morten Mjelde, Laurence Pilard, and Sébastien Tixeuil. A self-stabilizing 2/3-approximation algorithm for the maximum matching problem. In Sandeep S. Kulkarni and André Schiper, editors, *Stabilization, Safety, and Security of Distributed Systems, 10th International Symposium (SSS 2008)*, Lecture Notes in Computer Science, Detroit, November 2008. Springer-Verlag Berlin Heidelberg.

[35] Toshimitsu Masuzawa and Sébastien Tixeuil. Bounding the impact of unbounded attacks in stabilization. In Ajoy K. Datta and Maria Gradinariu, editors, *Eighth International Symposium on Stabilization, Safety, and Security on Distributed Systems (SSS 2006)*, Lecture Notes in Computer Science, pages 440–453, Dallas, Texas, November 2006. Springer Verlag.

[36] Toshimitsu Masuzawa and Sébastien Tixeuil. On bootstrapping topology knowledge in anonymous networks. In Ajoy K. Datta and Maria Gradinariu, editors, *Eighth International Symposium on Stabilization, Safety, and Security on Distributed Systems (SSS 2006)*, Lecture Notes in Computer Science, pages 454–468, Dallas, Texas, November 2006. Springer Verlag.

[37] Toshimitsu Masuzawa and Sébastien Tixeuil. Stabilizing link-coloration of arbitrary networks with unbounded byzantine faults. *International Journal of Principles and Applications of Information Science and Technology (PAIST)*, 1(1):1–13, December 2007.

[38] Toshimitsu Masuzawa and Sébastien Tixeuil. On bootstrapping topology knowledge in anonymous networks. *ACM Transactions on Adaptive and Autonomous Systems (TAAS)*, 2008.

[39] Toshimitsu Masuzawa and Sébastien Tixeuil. Quiescence of self-stabilizing gossiping among mobile agents in graphs. In *Proceedings of 15th International Colloquium on Structural Information and Communication Complexity (Sirocco 2008)*, Lecture Notes in Computer Science, Villars-sur-Ollon, Switzerland, June 2008. Springer-Verlag.

[40] Olivier Bournez Johanne Cohen Loubna Echabbi Sylvie Delaët, Dominique Barth. Existence of a Nash equilibria in a pricing game adapted to bgp. Technical report, LRI, 2006.

[41] Sébastien Tixeuil. *Vers l'Auto-stabilisation des Systèmes à Grande Echelle*. Hdr, Université Paris-Sud XI, Orsay, France, May 2006.

[42] Sébastien Tixeuil. *Wireless Ad Hoc and Sensor Networks*, chapter Fault-tolerant distributed algorithms for scalable systems. ISTE, October 2007. ISBN: 978 1 905209 86.

[43] Adnam Vora, Mikhail Nesterenko, Sébastien Tixeuil, and Sylvie Delaët. Universe detectors for sybil defense in ad hoc wireles networks. In *International Conference on Stabilization, Safety, and Security (SSS 2008)*, Lecture Notes in Computer Science. Springer-Verlag, November 2008.

C.4 La liste des éléments de valorisation

inventorie les items suivants : - brevets nationaux et internationaux, licences, et autres éléments de propriété intellectuelle consécutifs au projet. - logiciels et tout autre prototype - actions de normalisation - les conséquences du projet (lancement de produit ou service, nouveau projet, contrat,...) - le développement d'un nouveau partenariat, - la création d'une plate-forme à la disposition d'une communauté - autres (ouverture internationale,..) Elle en précise les partenariats éventuels. Dans le cas où des livrables ont été spécifiés dans l'annexe technique, on présentera ici un bilan de leur fourniture.

Le projet n'a pas mené directement à des brevets. Un certain nombre de logiciels académiques ont été développés. Le projet a surtout permis la création d'une collaboration durable, qui n'existait pas, entre les partenaires du projet.

Plusieurs soumissions ANR sont nées des suites de ce projet.

D. Rapport scientifique

Le rapport comporte deux parties : le résumé et le mémoire. Ce rapport doit être rédigé dans une forme adaptée à sa diffusion, en particulier via des archives ouvertes (http://hal.archives-ouvertes.fr/). Confidentialité-diffusion : les questions de confidentialité seront réglées en accord avec les unité support et le coordinateur du projet.

D.1. Le résumé consolidé

, maximum 1 page figures comprises, rappellera les objectifs du projet, les principaux résultats et conclusions et discutera de la concordance entre ceux-ci. On précisera aussi l'apport du projet aux équipes partenaires en insistant sur le différentiel lié à celui-ci. La situation dans la compétition internationale (état de l'art) sera précisée et on discutera l'éventuel effet de levier qu'il a pu activer, les retombées prévisibles et les perspectives ouvertes. Du fait de la possible diffusion de ce résumé auprès d'un large public, il ne fera pas mention de résultats confidentiels et utilisera un vocabulaire adapté mais n'excluant pas les termes techniques. Il en sera fourni une version française et une version en anglais.

Résumé en Français Dans le résumé qui suit, les numéros entre parenthèses désignent les tâches en rapport.

L'objectif du projet SOGEA était de contribuer à comprendre la théorie algorithmique des jeux et ses applications à la conception d'algorithmes distribués sûrs, avec des applications privilégiées tirées du routage interdomaine sur l'internet et des réseaux de capteurs distribués.

Nous nous sommes intéressés d'une part à modéliser le routage interdomaine dans l'Internet (1.1). Un modèle du routage interdomaine en présence de partenaires économiques a été complété par la prise en compte de différents aspects économiques (1.2, 2.1, 2.3). Nous avons aussi travaillé sur la modélisation, comme jeu, après discussions avec Alcatel, d'un problème de choix d'investissement en termes de nouvelles connexions dans un réseau interdomaine: résultats théoriques sur la complexité (3.1), caratérisation d'équilibres de Nash et utilisation de méthodes d'apprentissage stochastique pour identifier des équilibres mixtes (3.4).

Au niveau des aspects complexité (3.1, 3.2, 3.3), nous nous sommes concentrés sur l'approximation des Equilibres et sur la vérification approchée de systèmes probabilistes. Nous avons d'une part partiellement généralisé la vérification approchée introduite dans Fischer et al. aux systèmes probabilistes et obtenu partiellement des testeurs pour décider si la probabilité qu'un système vérifie une propriété est supérieure à un seuil λ.

D'autre part, nous avons considèré des jeux à information complète avec $r \geq 2$ joueurs, et étudié les Equilibres de Nash approchés dans le sens additif et multiplicatif lorsque le nombre de stratégies pures de chaque joueur est n. On sait qu'il n'existe pas de FPTAS pour ce problème et la principale question ouverte est de savoir s'il existe un PTAS. Nous avons prolongé les résultats connus en considérant un nombre arbitraire de joueurs et obtenu des bornes inférieures sur la taille des supports.

Nous avons par ailleurs introduit un modèle d'autostabilisation en présence de partenaires économiques aux intérêts divergents, que nous avons appelé autostabilisation égoiste (1.3). Ce modèle original a été étudié dans le cadre de différents algorithmes comme des algorithmes de construction d'arbres, en présence de plusieurs groupes de joueurs (2.2, 2.3, 2.4, 3.4) Nous avons relié la difficulté de ce problème aux classes de complexité classiques, et proposé plusieurs protocoles utilisables en pratique. Cette étude prolonge une étude préliminaire de l'autostabilisation en présence de partenaires byzantins (1.6, 2.4, 2.5).

Nous avons d'autre part travaillé sur la modélisation du dynamisme en théorie des jeux (1.4). Nous avons ainsi prouvé la convergence d'une dynamique d'apprentissage dans les problèmes de routage (1.5). Nous avons aussi contribué à comprendre la puissance du modèle des protocoles de population, qui a été introduit comme un modèle de réseaux de capteurs. Nous avons formalisé le modèle dans le cas de grandes populations, et nous avons caractérisé la puissance de calcul du modèle obtenu pour certains cas.

Summary in English In the following text, the numbers between parentheses correspond to related tasks.

The objective of the SOGEA project was to contribute to algorithmic game theory and its applications to the conception of safe distributed algorithms, with privilegied applications taken from interdomain routing and from sensor networks.

We first focused on modeling interdomain routing in Internet (1.1). A model of interdomain routing in presence of economic partners has been established, taking into account various economic aspects (1.2, 2.1, 2.3). We also worked on the modeling, as a game, after some discussions with Alcatel, of a problem of choice of investment of new connexions in an interdomain network: we obtained theoretical results (3.1), caracterisations of Nash equilibria, and use of stochastic learning techniques to identify mixed equilibria (3.4).

Concerning complexity aspects (3.1, 3.2, 3.4), we focused on the problem of approximating equilbria and on the approximated verification of probabilistic systems. On one hand, we partially generalized the approximate verification introduced by Fisher et al. to probabilistic systems, and we partially obtained some testers to decide whether the probability that a system satisfies a property is greater than some λ. On the other hand, we considered games with imperfect information with more than two players, and we studied Nash approximate equilibria in the additive and multiplicative sense when the number of pure strategies of each player is n. It is known that there is no FPTAS to this problem, but the main open problem is to understand whetheer there exists a PTAS. We extended known results by considering an arbitrary number of players, and obtained some lower bounds on the size of supports.

We also introduced a model of self-stabilization in presence of economic partners with diverging interests. This model has been called "selfish-self-stabilization" (1.3). This original model has been studied with various algorithms and problems like the problem of building spanning trees. The difficulty of the problem has been linked with classical complexity classes, and several algorithms that work in practise have been proposed. This work extends the preliminary work on self-stabilization with byzantine partners done in the project.

We also worked on the modeling of dynamism in game theory (1.4). We proved the convergence of a particular stochastic dynamic for routing problems (1.5). We also contributed to understand the power of protocol population models, that have been introduced as a model of sensor networks. We also formalized the model in the case of huge population, and we characterized the computational power of the obtained model, for some cases.

D.2. Le mémoire

maximum 5 pages, couvre la totalité de la durée du projet. Il doit être compris comme une synthèse auto-suffisante rappelant les objectifs, le travail réalisé et les résultats obtenus mis en perspective avec les attentes initiales et l'état de l'art. C'est un document d'un format semblable à celui des articles scientifiques ou des monographies. Il doit refléter le caractère collectif de l'effort fait par les partenaires au cours du projet. Le coordinateur prépare ce rapport sur la base des contributions de tous les partenaires. Une version préliminaire en est soumise à l'unité support dans un délai de deux mois suivant la fin du projet. La version finale, incorporant le cas échéant les commentaires et suggestions de l'unité support et la réponse des contractants, sera soumise en deux exemplaires-papier et un exemplaire électronique.

Format du mémoire On donne ici des indications sur le contenu possible du mémoire. En plus du résumé mentionné ci-dessus, celui-ci pourrait détailler : 1. Enjeux et problématique, état de l'art 2. Matériel et méthodes 3. Résultats 4. Discussion sur le degré de réalisation des objectifs initiaux, les verrous restant à franchir, les ruptures, les élargissements 5. Conclusions et recommandations sur l'exploitation et la dissémination des résultats. Ce mémoire peut être accompagné pour certains programmes de rapports annexes plus détaillé

Overall objectives Classical distributed algorithms have been mostly conceived under the hypothesis that each involved partner has no incentive in doing something different from what it is expected to do, without taking into account the possible economical interests of each partner.

This may be the source of serious security problems. Indeed, to guarantee for example the security of a protocol for telecommunication networks, or the security of a sensor-driven database, one must be able to guarantee that the presence of one or few adverse partners, could not lead to serious situations, or completely deteriorate the performances of the whole network. This becomes more and more crucial, as the size of the networks is increasing, or as the number of economical partners involved in today's networks is increasing.

Mathematical game theory studies the equilibria reached by rational players in competition, and Mechanism theory the possible games which may lead to equilibria. In recent years, algorithmic versions of these theories have been developed to study efficient methods to compute or verify approximate equilibria.

Algorithms from distributed algorithm theory must be reconsidered under these points of view, in order to guarantee their correct behavior and performances even in presence of partners with divergent economical interests, or that could be subject to some selfish behaviors.

The objective of the SOGEA project is contribute to understand algorithmic game theory and its applications to classical and self-stabilizing distributed algorithms.

By this action, we proposed specifically the following points:

- produce some models and notions of stability from game theory adapted to problems related to security of distributed systems;

- produce some distributed algorithms, preferably self-stabilizing, for simple models coming from problems related to inter-domain routing in telecommunication networks, and sensor networks with guarantees of security in terms of these notions of stability;
- the proposed solutions are validated analytically with a well understood underlying complexity theory.

Models of Interdomain Routing We presented both a game theoretic and a distributed algorithmic approach for the transit price negotiation problem in the interdomain routing framework in several papers.

In a first model, we dealt with inter-domain routing management from an economical point of view. We presented a game theory based costing model that maps BGP peers (autonomous systems belonging to different operators) into a strategic (selfish) agents competing for transit traffic as a service provided and charged to their peers. Indeed, in our model each operator fixes a price to each neighbor for each transit traffic unit. Then, BGP routing choice is made based on a minimum cost criterion where the goal of each operator is to minimize its costs. We investigate some particular strategies of updating prices that operators can use locally in order to minimize their costs. We focus on BGP stabilization properties related to such strategies from a simulation point of view.

We proved in [40] that this model does not always admit pure equilibria, even for a very few number of nodes.

The analysis of the centralized transit price negotiation problem shows that the only one non cooperative equilibrium is when the lowest cost provider takes all the market. The perspective of the game being repeated makes cooperation possible while maintaining higher prices. We considered then the system under a realistic distributed framework and simulate its behaviour under a simple price adjustment strategy and analyse whether it matches the theoretical results in [5].

We presented a distributed learning algorithm for optimising transit prices in a negotiation problem in the interdomain routing framework. We presented a combined game theoretic and a distributed algorithmic analysis where the notion of Nash equilibrium with the first approach model meets the notion of stability in the second. We show that minimum cost providers can learn how to set their prices strategically according to a Nash equilibrium. We validated our theoretic model by simulations that confirm the expected outcome. Moreover, we show that players that deviate from the proposed rule in order to change the outcome of the game seem to not affect the expected dynamic of the system in [21].

Dynamical Aspects A discussion on dynamical aspects and their modeling in game theory has been done in [9].

Population protocols have been introduced as a model of sensor networks consisting of very limited mobile agents with no control over their own movement. A population protocol corresponds to a collection of anonymous agents, modeled by finite automata, that interact with one another to carry out computations, by updating their states, using some rules. Their computational power has been investigated under several hypotheses but always when restricted to finite size populations. In particular, predicates stably computable in the original model have been characterized as those definable in Presburger arithmetic. We studied mathematically the convergence of population protocols when the size of the population goes to infinity in [11]. We did so by giving general results, that we illustrate through the example of a particular population protocol for which we even obtain an asymptotic development. This example shows in particular that these protocols seem to have a rather different computational power when a huge population hypothesis is considered.

Complexity Aspects We considered in [32] games of complete information with $r \geq 2$ players, and studied approximate Nash equilibria in the additive and multiplicative sense, where the number of pure strategies of the players is n. We establish a lower bound on the size of the support of strategy profiles which achieve an ϵ-approximate equilibrium, for $\epsilon < \frac{r-1}{r}$ in the additive case, and $\epsilon < r-1$ in the multiplicative case. We exhibit polynomial time algorithms for additive approximation which respectively compute an $\frac{r-1}{r}$-approximate equilibrium with support sizes at most 2, and which extend the algorithms for 2 players with better than 1/2-approximations to compute ϵ- equilibria with $\epsilon < \frac{r-1}{r}$. Finally, we investigate the sampling based technique for computing approximate equilibria of Lipton et al.[12] with a new analysis, that instead of Hoeffding's bound uses the more general McDiarmid's inequality. In the additive case we show that for $0 < \epsilon < 1$, an ϵ-approximate Nash equilibrium with support size $\frac{2rln(nr+r)}{\epsilon^2}$ can be obtained, improving by a factor of r the support size of [12]. We derive an analogous result in the multiplicative case where the support size depends also quadratically on $g-1$, for any lower bound on the payoffs of the players at some given Nash equilibrium.

We introduced approximate data exchange in [22], by relaxing classical data exchange problems such as Consistency and Typechecking to their approximate versions based on Property Testing. This problem can be seen as a game between a source and a destination that must send some information, fullfilling some imposed schemas. It provides a natural framework for consistency and safety questions, which first considers approximate solutions and then exact solutions obtained with a Corrector. We consider a model based on transducers of words and trees, and study ϵ-Consistency, i.e., the problem of deciding whether a given source instance I is ϵ-close to a source I', whose image by a transducer is also ϵ-close to a target schema. We prove that ϵ-Consistency has an ϵ-tester, i.e. can be solved by looking at a constant fraction of the input I. We also show that ϵ-Typechecking on words can be solved in polynomial time, whereas the exact problem is PSPACE-complete. Moreover, data exchange settings can be composed when they are close.

Inspired by property testing, we relaxed in [29] the classical satisfiability $U \vdash F$ between a finite structure U of a class K and a formula F, to a notion of ϵ-satisfiability $U \vdash F$, and the classical equivalence $F1 \equiv F2$ between two formulas F1 and F2, to ϵ-equivalence $F1 \equiv_\epsilon F2$ for $\epsilon > 0$. We considered the class of strings and trees with the edit distance with moves, and show that these approximate notions can be efficiently decided. We give a tester for equality and membership in any regular language, in time independent of the size of the structure. Using our geometrical embedding, we can also test the equivalence between two regular properties on words, defined by monadic second order formulas. Our equivalence tester has polynomial time complexity in the size of the automaton (or regular expression), for a fixed ϵ, whereas the exact version of the equivalence problem is PSPACE-complete. Last, we extended the geometric embedding, and hence the tester algorithms, to infinite regular languages and to context-free languages. For context-free languages, the equivalence tester has an exponential time complexity, whereas the exact version is undecidable.

Self-Stabilization with Adversaries As a new challenge of containing the unbounded influence of Byzantine processes in self-stabilizing protocols, [35] introduces a novel concept of *strong stabilization*. The strong stabilization relaxes the requirement of *strict stabilization* so that processes beyond the containment radius are allowed to be disturbed by Byzantine processes, but only a limited number of times. A self-stabilizing protocol is (t, c, f)-strongly stabilizing if any process more than c hops away from any Byzantine process is disturbed at most t times in a distributed system with at most f Byzantine processes. Here c denotes the *containment radius* and t denotes the *containment times*. The possibility and the effectiveness of the strong stabilization is demonstrated using *tree orientation*. It is known that the tree orientation has no strictly stabilizing protocol with a constant containment radius. This paper first shows that the problem has no constant bound of the containment radius in a tree with two Byzantine processes even when we allow processes beyond the containment radius to be disturbed any finite number of times. Then we consider the case of a single Byzantine process and present a $(1, 0, 1)$-strongly stabilizing protocol, which achieves optimality in both containment radius and times.

We generalized the classic dining philosophers problem to separate the conflict and communication neighbors of each process in [19]. Communication neighbors may directly exchange information while conflict neighbors compete for the access to the exclusive critical section of code. This generalization is motivated by a number of practical problems in distributed systems including problems in wireless sensor networks. We present a self-stabilizing deterministic algorithm — KDP that solves a restricted version of the generalized problem where the conflict set for each process is limited to its k-hop neighborhood. Our algorithm is terminating. We formally prove KDP correct and evaluate its performance. We then extended KDP to handle fully generalized problem. We further extend it to handle a similarly generalized drinking philosophers problem. We describe how KDP can be implemented in wireless sensor networks and demonstrate that this implementation does not jeopardize its correctness or termination properties.

In [38], we quantified the amount of "practical information (*i.e.* views obtained from the neighbors, colors attributed to the nodes and links) to obtain "theoretical information (*i.e.* the local topology of the network up to distance k) in anonymous networks. In more details, we show that a coloring at distance $2k + 1$ is necessary and sufficient to obtain the local topology at distance k that includes outgoing links. This bound drops to $2k$ when outgoing links are not needed. A second contribution of this paper deals with color bootstrapping (from which local topology can be obtained using the aforementioned mechanisms). On the negative side, we show that *(i)* with a distributed daemon, it is impossible to achieve deterministic color bootstrap, even if the whole network topology can be instantaneously obtained, and *(ii)* with a central daemon, it is impossible to achieve distance m when instantaneous topology knowledge is limited to $m - 1$. On the positive side, we show that *(i)* under the k-central daemon, deterministic self-stabilizing bootstrap of colors up to distance k is possible provided that k-local topology can be instantaneously obtained, and *(ii)* under the distributed daemon, probabilistic self-stabilizing bootstrap is possible for any range.

The Sybil attack in unknown port networks such as wireless is not considered tractable. A wireless node is not capable of independently differentiating the universe of real nodes from the universe of arbitrary non-existent fictitious nodes created by the attacker. Similar to failure detectors, we proposed in [43] to use *universe detectors* to help nodes determine which universe is real. In this paper, we (i) define several variants of the neighborhood discovery problem under Sybil attack (ii) propose a set of matching universe detectors (iii) demonstrate the necessity of additional topological constraints for the problems to be solvable: node density and communication range; (iv) present $\mathcal{SAND}$ — an algorithm that solves these problems with the help of appropriate universe detectors, this solution demonstrates that the proposed universe detectors are the weakest detectors possible for each problem.

The matching problem asks for a large set of disjoint edges in a graph. It is a problem that has received considerable attention both from the sequential and the self-stabilizing communities. Previous work has resulted in self-stabilizing algorithms for computing a maximal (1/2-approximation) matching in a general graph, as well as computing a 2/3-approximation on more specific graph types. In [34], we proposed the first self-stabilizing algorithm for finding a 2/3-approximation to the maximum matching problem in a general graph. We show that our new algorithm stabilizes in at most exponential time under a distributed adversarial daemon, and $O(n^2)$ rounds under a distributed fair daemon, where n is the number of nodes in the graph.

In [25, 23], we tackled the open problem of snap-stabilization in message-passing systems. Snap-stabilization is a nice approach to design protocols that withstand transient faults. Compared to the well-known self-stabilizing approach, snap-stabilization guarantees that the effect of faults is contained immediately after faults cease to occur. Our contribution is twofold: we show that (1) snap-stabilization is impossible for a wide class of problems if we consider networks with finite yet unbounded channel capacity; (2) snap-stabilization becomes possible in the same setting if we assume bounded-capacity channels. We propose three snap-stabilizing protocols working in fully-connected networks. Our work opens exciting new research perspectives, as it enables the snap-stabilizing paradigm to be implemented in actual networks.

In [39], we considered gossiping among mobile agents in graphs: agents move on the graph and have to disseminate their initial information to every other agent. We focus on self-stabilizing solutions for the gossiping problem, where agents may start from arbitrary locations in arbitrary states. Self-stabilization requires (some of the) participating agents to keep moving forever, hinting at maximizing the number of agents that could be allowed to stop moving eventually. This paper formalizes the self-stabilizing agent gossip problem, introduces the quiescence number (i.e., the maximum number of eventually stopping agents) of self-stabilizing solutions and investigates the quiescence number with respect to several assumptions related to synchrony, whiteboard availability, and anonymity.

Properly locating sensor nodes is an important building block for a large subset of wireless sensor networks (WSN) applications. As a result, the performance of the WSN degrades significantly when misbehaving nodes report false location and distance information in order to fake their actual location. In [26], we proposed a general distributed deterministic protocol for accurate identification of faking sensors in a WSN. Our scheme does *not* rely on a subset of *trusted* nodes that are not allowed to misbehave and are known to every node in the network. Thus, any subset of nodes is allowed to try faking its position. As in previous approaches, our protocol is based on distance evaluation techniques developed for WSN. On the positive side, we show that when the received signal strength (RSS) technique is used, our protocol handles at most $\lfloor \frac{n}{2} \rfloor - 2$ faking sensors. Also, when the time of flight (ToF) technique is used, our protocol manages at most $\lfloor \frac{n}{2} \rfloor - 3$ misbehaving sensors. On the negative side, we prove that no deterministic protocol can identify faking sensors if their number is $\lceil \frac{n}{2} \rceil - 1$. Thus our scheme is almost optimal with respect to the number of faking sensors. We discuss application of our technique in the trusted sensor model. More precisely our results can be used to minimize the number of trusted sensors that are needed to defeat faking ones.

We wrote a survey in [42].

Selfish Self-Stabilization Stabilizing distributed systems expect all the component processes to run predefined programs that are externally mandated. In Internet scale systems, this is unrealistic, since each process may have selfish interests and motives related to maximizing its own payoff. We formulated the problem of selfish stabilization that shows how competition blends with cooperation in a stabilizing environment in [20].

This has been intensively extended in [16]. This paper addresses a specific problem in this domain. Given a graph G = (V , E), assume that there are p different subsets (or classes or colors) of nodes. For each subset or color, there is a separate cost function that maps the set of edges to the set of positive integers. Starting from an arbitrary initial configuration, the p different classes of nodes cooperate with one another to form a spanning tree rooted at a designated node, and at the same time compete against each other to minimize their cost of communication with the root node. The communication cost may depend on various factors: for example, ownership of the routers may be a factor in determining the cost of routing traffic for any class of nodes. The processes are free to choose a strategy from a given set of strategies, and switch strategy to satisfy their individual needs. We examined strategies under which, starting from an arbitrary initial configuration, these classes of processes can stabilize to an equilibrium configuration after which no process can unilaterally decrease its communication cost to the root. We showed that with three (or more) different classes of nodes, there exists the possibility that no equilibrium is possible, and proves that determining whether a particular setting admits an equilibrium is NP- complete. On the positive side, we presented evidence that when there are two (or less) classes, an equilibrium always exists. A (weakly) stabilizing distributed algorithm is provided to construct an equilibrium tree. Alternative strategies for participating nodes are also discussed.