



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*Stability of a local greedy distributed routing
algorithm*

Florian Huc — Christelle Molle — Nicolas Nisse — Stephane Perennes — Herve Rivano

N° 6871

March 2009

Thème COM

A large, light gray stylized letter 'R' is positioned to the left of the text. A horizontal gray brushstroke is located below the text.

*R*apport
de recherche



Stability of a local greedy distributed routing algorithm

Florian Huc^{*†} , Christelle Molle^{‡§} , Nicolas Nisse^{‡§} , Stephane Perennes^{‡§} ,
Herve Rivano^{‡§}

Thème COM — Systèmes communicants
Projets Mascotte

Rapport de recherche n° 6871 — March 2009 — 11 pages

Abstract: In this work, we study the problem of routing packets between undifferentiated sources and sinks in a network modeled by a multigraph. We provide a distributed and local algorithm that transmits packets hop by hop in the network and study its behaviour. At each step, a node transmits its queued packets to its neighbours in order to optimize a local gradient. This protocol is thus greedy since it does not require to record the history about the past actions, and lazy since it only needs informations of the neighborhood.

We prove that this protocol is however optimal in the sense that the number of packets stored in the network stays bounded as soon as the sources injects a flow that another method could have exhausted. We therefore reinforce a result from the literature that worked for differentiated suboptimal flows.

Key-words: distributed algorithm, greedy, stability

This work has been partly funded by the European project IST-FET IP AEOLUS.

* TCS-Sensor Lab, Centre Universitaire d'Informatique, Université de Genève, Suisse.

† florian.huc@gmail.com

‡ MASCOTTE, INRIA, I3S, CNRS, Univ. Nice Sophia, Sophia Antipolis, France.

§ firstname.lastname@sophia.inria.fr

Stabilité d'un algorithme distribué de routage de type gradient glouton local

Résumé : Dans cet article, nous nous intéressons au problème du routage depuis des sources vers des puits indifférenciés dans un réseau modélisé par un multigraphe. Nous étudions le comportement d'un algorithme distribué et local de transmission de paquets de proche en proche dans le réseau. A chaque étape, un nœud transmet les paquets qu'il a en transit vers ses voisins de manière à optimiser un gradient local. Ce protocole est ainsi glouton puisque ne prenant pas en compte l'historique du réseau, et "feignant" puisque ne considérant que l'information de ses voisins.

Nous montrons que ce protocole est néanmoins optimal dans le sens où le nombre de paquets en transit dans le réseau reste borné tant que les sources injectent un flot qu'une autre méthode saurait écouler. Nous renforçons donc, sous nos hypothèses, un résultat de la littérature valable pour des flots différenciés sous-optimaux.

Mots-clés : algorithme distribué, glouton, stabilité

1 Introduction

The actual progress of networks involves an increasing interest for distributed algorithms that use only few information about the network. We study a local protocol for routing packets dynamically. We prove that the protocol is *stable*, i.e. the number of packets stored at the nodes of the network is bounded (does not grow to infinity).

In previous works, Srikant et al. [4] studied distributed and localized algorithms to transmit packets in a network. In their study, they do not deal with the routing and only focus on the call scheduling for one-hop communications when calls are matching and packets enter the network continuously. They base their work on an article by Tassioulas et al. [2] who have proposed a family of stable algorithms. In both of these cases, packets are injected into the network following a stochastic process that respects a *strict* feasibility constraint, saying that the number of added packets at a time is always strictly lower than the value of the maximum flow.

Other works have considered processes in which packets are given by an adversary who wants to make the protocol fail [3].

In this work, we consider a simplified network model in which sources inject packets into the network, then lazy nodes forward these packets according to a local greedy gradient computation with the only information of their neighbours' state, and sinks extract the packets from the network. This behaviour can be related to the distributed algorithm for the maximum flow problem proposed by Goldberg et Tarjan [1].

We show that, in the case with undifferentiated sources and sinks, our protocol is optimal in the sense that the number of packets stored in the network stays bounded as soon as the sources injects a flow that another method could have exhausted. In particular, our protocol stays stable without the strict feasibility condition.

1.1 The network model

Let $G = (V, E)$ be a multigraph modeling the considered network. We denote Δ the maximum degree of G , and $\Gamma(u)$ the neighborhood of node $u \in V$.

To each vertex is associated a queue state which represents the number of packets waiting to be transmitted at this node. We represent this queue state by $q_t(v)$ for $v \in V$ and a given time step t , also called the *height* of v . Let $\mathcal{S} \subseteq V$ and $\mathcal{D} \subseteq V$ be respectively the sets of source and destination nodes (Fig. 1).

The network is synchronous and at each time step:

- each source $s \in \mathcal{S}$ injects $in(s)$ packets in its queue,
- each link can transmit at most 1 packet, and this packet can be lost without any information at the sending node,
- each sink $d \in \mathcal{D}$ extracts $out(d)$ packets of its queue.

All links can eventually transmit at the same time, and the set of links that simultaneously transmit at time t is denoted E_t .




Figure 1: The multigraph G representing the network.

1.2 The greedy and lazy nodes

The network nodes run Algorithm 1 simultaneously. They only need to get access to the queue state of their neighbours.

Algorithm 1: Algorithm *LGG* : Local greedy gradient.

```

 $E_t(u) \leftarrow \emptyset$ 
 $q \leftarrow q_t(u)$ 
 $list(u) \leftarrow$  ordonner  $\Gamma(u)$  par  $q_t$  croissant
for all  $v \in list(u)$  do
  if  $q_t(u) > q_t(v)$  &&  $q > 0$  then
     $E_t(u) \leftarrow E_t(u) \cup \{(u, v)\}$ 
     $q \leftarrow q - 1$ 

```

At each time step t , each source s injects $in(s)$ packets in its queue. Then, each node u transmits 1 packet on each of its outgoing arcs with destination v that has the smallest height, as soon as u still has packets in its queue. In particular, if u would have sent more than $q_t(u)$ packets, then it chooses to send to its $q_t(u)$ neighbours of smallest height. This choice actually has no impact on the system stability. The set of transmissions of u at time t is denoted $E_t(u)$, and $\cup_{u \in V} E_t(u) = E_t$.

Packets destined to node v are deleted from u 's queue, and, for each successful transmission, 1 packet is added to v 's queue.

Finally, each sink d remove $\min\{out(d), q_t(d)\}$ packets from its queue and step t is over.

2 Study of stability

Let G^* be the multigraph obtained from G by adding a virtual source s^* and a virtual sink d^* , with $in(s)$ links between s^* and s for all $s \in \mathcal{S}$, and $out(d)$ links between d and d^* for all $d \in \mathcal{D}$.

In the following, we denote f^* the value of a s^* - d^* maximum flow in G^* .

Definition 1 (Feasible arrival rate) An arrival rate is feasible if $\sum_{s \in \mathcal{S}} in(s) \leq f^*$.

Theorem 1 If the arrival rate is feasible, then the system is stable. Otherwise, the number of packets stored in the network may diverge with the time no matter what algorithm is used.

At each time step t , the *network state* is measured by $P_t = \sum_{u \in V} q_t^2(u)$. We first remark that the system can diverge if $\sum_{s \in \mathcal{S}} in(s) > f^*$. Indeed, without any additional assumptions on the packet loss, we can assume that all packets are delivered. It is then enough to look at a \mathcal{S} - \mathcal{D} minimum cut (A, B) (of value f^*), with $\mathcal{S} \subseteq A$. At each step at most f^* packets leave A whereas $\sum_{s \in \mathcal{S}} in(s) (> f^*)$ enter it. So P_t strictly increases at each step.

In order to prove Theorem 1, we investigate three cases depending on the value of the difference between the maximum flow and the arrival rate. We say that the network is *unsaturated* if $\sum_{s \in \mathcal{S}} in(s) < f^*$, and *saturated* if $\sum_{s \in \mathcal{S}} in(s) = f^*$. Moreover, the network is *uniformly unsaturated* if it exists a flow Φ from s^* to d^* such that, for each source $s \in \mathcal{S}$, $in(s) < \Phi(s)$ units of flow pass through s at each step. Otherwise, it exists at least one *saturated* source $s' \in \mathcal{S}$ such that $in(s') = \Phi(s')$ and we call this case *non uniformly unsaturated*.

In the next section, we prove the uniformly unsaturated case, covering a part of Theorem 1 when the arrival rate is strictly feasible. We then introduce a more general node representation for the general proof of Theorem 1 (including the saturated case).

2.1 Proof of the uniformly unsaturated case

By definition, it exists a flow Φ from s^* to d^* such that, for all source $s \in \mathcal{S}$, $\Phi(s) > in(s)$.

Let us consider the evolution of the network state between step t and $t + 1$:

$$P_{t+1} = \sum_{v \in V} q_{t+1}^2(v) = \sum_{u \in V} q_t^2(u) + \sum_{v \in V} (q_{t+1}(v) - q_t(v))^2 + 2 \sum_{v \in V} q_t(v)(q_{t+1}(v) - q_t(v)) \quad (1)$$

For all $v \in V$, $(q_{t+1}(v) - q_t(v)) \leq \Delta$, then:

$$P_{t+1} \leq P_t + \delta_t + n\Delta^2, \text{ with } \delta_t = \sum_{v \in V} q_t(v)(q_{t+1}(v) - q_t(v)).$$

Equivalently, δ_t can be defined in function of the links in E_t used by *LGG* at time t for the transmissions. In the following, $e = (u, v) \in E_t$ is oriented to indicate that the packet goes from u to v . Then, δ_t can be formulated as follows:

$$\delta_t = \sum_{s \in \mathcal{S}} q_t(s)in(s) - \sum_{d \in \mathcal{D}} q_t(d) \min\{out(d), q_t(d)\} + \sum_{(u,v) \in E_t} (q_t(v) - q_t(u)) \quad (2)$$

We now compare the variation of P_t during an execution of *LGG* to the one obtained by pushing the packets along the paths allowing a maximum flow. Let us consider the set of f^* edge-disjoint paths between the sources and the sinks used by flow Φ , and E_t^Φ the set of links (source-to-destination oriented) of these paths selected at time t . By summing the difference of the potential on each hop along these paths, we get:

$$\sum_{(u,v) \in E_t^\Phi} (q_t(v) - q_t(u)) - \sum_{d \in \mathcal{D}} q_t(d) \min\{out(d), q_t(d)\} = - \sum_{s \in \mathcal{S}} q_t(s)\Phi(s).$$

Let us now study the sum of the difference of the potential on the links used by *LGG*:

$$\begin{aligned} \sum_{(u,v) \in E_t} (q_t(v) - q_t(u)) &= \sum_{(u,v) \in E_t^\Phi} (q_t(v) - q_t(u)) - \sum_{(u,v) \in E_t^\Phi \setminus E_t} (q_t(v) - q_t(u)) \\ &\quad + \sum_{(u,v) \in E_t \setminus E_t^\Phi} (q_t(v) - q_t(u)). \end{aligned}$$

By definition of *LGG*, for all $e = (u, v) \in E_t$, $q_t(v) - q_t(u) < 0$. So, $\sum_{(u,v) \in E_t \setminus E_t^\Phi} (q_t(v) - q_t(u)) < 0$. Moreover if $e = (u, v) \in E_t^\Phi \setminus E_t$, then, again by definition of *LGG*, either $q_t(v) \geq q_t(u)$ or $q_t(u) \leq \Delta$. Indeed if $q_t(v) < q_t(u)$, our algorithm must send 1 packet from u to v , unless u has already sent all its available packets in $q_t(u)$. So $\sum_{(u,v) \in E_t^\Phi \setminus E_t} (q_t(v) - q_t(u)) \geq \sum_{(u,v) \in E_t^\Phi \setminus E_t} (-\Delta) \geq -n\Delta^2$ and $\sum_{(u,v) \in E_t} (q_t(v) - q_t(u)) - \sum_{d \in \mathcal{D}} q_t(d) \min\{out(d), q_t(d)\} \leq -\sum_{s \in \mathcal{S}} q_t(s)\Phi(s) + n\Delta^2$.

From equation (2), we deduce that $\delta_t \leq \sum_{s \in \mathcal{S}} q_t(s)(in(s) - \Phi(s)) + n\Delta^2$. By definition of Φ , for all $s \in \mathcal{S}$, $in(s) < \Phi(s)$. So $\delta_t \leq n\Delta^2$ and, from equation 1, we can upper bound the evolution of the network state:

$$P_{t+1} - P_t \leq 2n\Delta^2. \quad (3)$$

We now show that: if $P_t \geq 4(f^* + 1)n^2\Delta^2$, then:

$$P_t - P_{t+1} > 2n\Delta^2. \quad (4)$$

Let us first assume that it exists a source $s \in \mathcal{S}$ such that $q_t(s) \geq 4n\Delta^2$. Then $\delta_t \leq -q_t(s) + n\Delta^2 < -3n\Delta^2$ and the result 4 is valid.

Secondly, we are now in the case where $q_t(s) < 4n\Delta^2$ for all $s \in \mathcal{S}$. If $P_t \geq 4(f^* + 1)n^2\Delta^2$, then it exists $v \in V \setminus \mathcal{S}$ such that $q_t(v) \geq 4(f^* + 1)n\Delta^2$. Let $v = u_1, u_2, \dots, u_k = d$ be a path from v to $d \in \mathcal{D}$ (eventually $d = v$). Then $\sum_{i < k, q_t(u_i) > q_t(u_{i+1})} (q_t(u_{i+1}) - q_t(u_i)) - q_t(u_k) \min\{out(u_k), q_t(u_k)\} \leq -q_t(v)$. This sum contributes negatively to $\sum_{(u,v) \in E_t} (q_t(v) - q_t(u)) - \sum_{d \in \mathcal{D}} q_t(d) \min\{out(d), q_t(d)\}$. From equation 2, we thus obtain that:

$$\delta_t \leq \sum_{s \in \mathcal{S}} q_t(s)in(s) - q_t(v) < f^* \cdot \max_{s \in \mathcal{S}} q_t(s) - q_t(v) \leq f^* \cdot 4n\Delta^2 - q_t(v) \leq -3n\Delta^2$$

and equation 4 is verified.

From 3 and 4 we deduce that, for all t , $P_t \leq 4(f^* + 1)n^2\Delta^2 + 2n\Delta^2$ which bounds the number of packets stored in the network at each time step and prove the strict stability of our algorithm. We remark that the packet losses here only improve the protocol stability.

3 Proof by induction on more general settings

The general proof of Theorem 1 is done by induction on the network size. The purpose of the induction is to prove that it exists, for all graph G , a constant $R(G)$ such that, for all

$R \geq R(G)$, our protocol is stable in G containing sources and destinations whose behaviour can be generalized as follows:

Definition 2 (Generalized sources) A generalized source s injects at most $in(s)$ packets in its queue at the beginning of each step.

Definition 3 (R -generalized destinations) A generalized destination d extracts at most $out(d)$ packets of its queue at the end of each step, and, given a constant of retention R :

- (i) if $q_t(d) > R$, then d extracts at most $\min\{out(d), q_t(d) - R\}$,
- (ii) for each $u \in \Gamma(d)$, d reveals a queue size $q'_t(d)$ defined as follows:
 - if $q_t(d) > R$, then d declare $q'_t(d) = q_t(d)$,
 - if $q_t(d) \leq R$, then d declare an height $q'_t(d) \leq R$.

Remark 1 The sources and sinks as defined in Section 1.1 are actually included in these new definitions of generalized sources and destinations. It is obvious to see that a "normal" source is a specific case of a generalized one when it injects exactly $in(s)$ packets at each time. And a "normal" destination is a R -generalized destination with $R = 0$.

3.1 Uniformly unsaturated case

If the network is uniformly unsaturated, then using Section 2.1 we prove that the network with generalized sources and destinations is stable, even if components A and B are respectively reduced to \mathcal{S} and \mathcal{D} (since $\sum_{s \in \mathcal{S}} in(s) < f^*$).

So we slightly change the proof of Section 2.1 to adapt to the generalized sources and sinks. The first change occurs in equation (2) that becomes:

$$\delta_t = \sum_{s \in \mathcal{S}} q_t(s) in(s) - \sum_{d \in \mathcal{D}} q_t(d) \min\{out(d), q_t(d) - R\} + \sum_{(u,v) \in E_t} (q_t(v) - q_t(u)) \quad (5)$$

But this has no impact on the proof of the upper bound $P_{t+1} - P_t \leq 2n\Delta^2$.

Then, in the proof of the lower bound, we have to modify the analysis when $q_t(s) < 4n\Delta^2$, for all $s \in \mathcal{S}$. If $P_t \geq 4(f^* + 1)n^2\Delta^2 + Rn$, then it exists $v \in V \setminus \mathcal{S}$ such that $q_t(v) \geq 4(f^* + 1)n\Delta^2 + R$. Let $v = u_1, u_2, \dots, u_k = d$ be a path from v to $d \in \mathcal{D}$ (eventually $d = v$). If $q_t(d) > R$, then d declare its real queue size and the rest of the proof does not change. If $q_t(d) \leq R$, then $\sum_{i < k, q_t(u_i) > q_t(u_{i+1})} (q_t(u_{i+1}) - q_t(u_i)) - q_t(u_k) \min\{out(u_k), q_t(u_k) - R\} \leq -q_t(v) + R$. So now:

$$\delta_t \leq \sum_{s \in \mathcal{S}} q_t(s) in(s) - q_t(v) + R < f^* \cdot \max_{s \in \mathcal{S}} q_t(s) - q_t(v) + R \leq f^* \cdot 4n\Delta^2 - q_t(v) + R \leq -3n\Delta^2$$

and the proof is done.

3.2 Non uniformly unsaturated case

If the network is non uniformly unsaturated, then it exists at least one source $s \in \mathcal{S}$ such that $in(s) = \Phi(s)$ for a feasible flow Φ . Let \mathcal{S}_s be the subset of \mathcal{S} maximal for inclusion containing all the saturated sources s' , such that $\Phi(s') = in(s')$. We look at the minimum cut (A', B') of value $\sum_{s' \in \mathcal{S}_s} in(s')$ that exists between \mathcal{S}_s and the destinations that is closest to the destinations (cf Fig. 2).




Figure 2: Minimum cut (A', B') such that the saturated sources are in A' .

Let us globally denote $\mathcal{S}(A')$ the set of sources in A' (i.e. the saturated sources \mathcal{S}_s), and $\mathcal{S}(B')$ the set of sources in B' ($\mathcal{S}(B') = \mathcal{S} \setminus \mathcal{S}(A')$). We can remark that A' acts as a single node for the protocol in B' , corresponding to a single generalized source s^* with arrival rate $in(s^*) = \sum_{s \in \mathcal{S}(A')} in(s)$.

We now look at the set $\mathcal{S}' \subseteq B'$ of nodes adjacent to the edges of the cut, i.e. containing the neighbours of A' in B' . We then have to deal with two cases whether destinations are in \mathcal{S}' or not.



3.2.1 $\mathcal{S}' \cap \mathcal{D} = \emptyset$

We can consider \mathcal{S}' as generalized sources in B' with arrival rate equal to their degree with A' represented by a single node s^* . By induction on the smaller network B' with generalized sources $\mathcal{S}(B') \cup \mathcal{S}'$ and R -generalized destinations \mathcal{D} , we get stability.

3.2.2 $\mathcal{S}' \cap \mathcal{D} \neq \emptyset$

The only changes occur at the generalized destinations that are linked to s^* . If a destination is linked to a source, then we can do the algebraic sum to delete the source and deal with a modified destination. For all $d \in \mathcal{S}' \cap \mathcal{D}$, let $out'(d) = out(d) + deg(s^*, d)$, where $deg(s^*, d) = |\{\text{the number of links between } s^* \text{ and } d\}|$. Then d acts as a R' -generalized destination in B' , with $R' = R - deg(s^*, d)$.

So B' verifies now the uniformly unsaturated condition and, by Section 3.1, has a bounded number of stored packets.


 Figure 3: Definition of a \mathcal{S} - \mathcal{D} cut in the network.

 Figure 4: Definition of a \mathcal{S} - \mathcal{D} cut and associated subsets \mathcal{S}' and \mathcal{D}' .

3.3 Saturated case

We consider a \mathcal{S} - \mathcal{D} minimum cut (A, B) , $\mathcal{S} \subseteq A$ and $\mathcal{D} \subseteq B$, and treat separately each of the two components (cf., Figure 4).

Different cases have to be considered whether $A = \mathcal{S}$, $B = \mathcal{D}$, or not. Let us denote $\mathcal{S}' \subseteq B$ (respectively $\mathcal{D}' \subseteq A$), the subset of nodes adjacent to the edges of the cut (see Fig. ??).

We first prove that component B is bounded.

3.3.1 $\mathcal{S}' \cap \mathcal{D} = \emptyset$

Since the cut can deliver at most 1 packet per edge to B , if a node $s' \in \mathcal{S}'$ is linked to k edges to A , then s' can deliver at most k packets to B and it may lose some packets by sending them to A . So we can consider this node as a generalized source with arrival rate $in(s') = |\{(v, s'), v \in A\}|$.

Looking at B , we can now view it as a smaller network with $|\mathcal{S}'|$ generalized sources. We have in B a smaller flow problem with $\sum_{s \in \mathcal{S}'} in(s) = f^*$. Packets may disappear from B (going to A), and a generalized source may create less packets than its maximal rate. From induction on the cardinal of $V(G)$, we conclude that the network B is stable.

3.3.2 $S' \cap \mathcal{D} \neq \emptyset$

As in the Section 3.2, we make the algebraic sum between the sources linked directly to the destinations. For each $d \in S' \cap \mathcal{D}$, if $out(d) - |\{(v, d), v \in A\}| \geq 0$, then we consider it as a R' -generalized destination with $out'(d) = out(d) - |\{(v, d), v \in A\}|$, and $R' = R - |\{(v, d), v \in A\}|$.

Otherwise, we consider it as a generalized source with $in(s) = |\{(v, d), v \in A\}| - out(d)$.

Considering this generalized sources and destinations, we obtain have a feasible arrival rate (recall we consider a minimal cut between the sources and the destinations, and that we removed as much capacity from the sources than from the input). Hence, by induction, this part of the network is stable.

We have proved that B has a bounded number of stored packets, so now we look at A and remark that B acts as a single node p for the protocol in A . We thus contract B into a node p such that $o(p) = \sum_{d \in B} out(d)$, and obtain a R_p -generalized destination verifying the definition 3 with $R_p = |B| \cdot M$, where M is a constant bounding the queue size of the nodes in B .

Remark 2 *If $|B| > 1$, we can directly apply the induction on the smaller network obtained to obtained that A has a bounded number of stored packets. Hence the following concerns the case $|B| = 1$.*

3.3.3 $S \cap \mathcal{D}' = \emptyset$

We know that B is bounded, so if a node in A adjacent to B is high enough, it will send packets to B . B can fireback packets only to nodes in A that have few packets. This is exactly the definition of a generalized destination. More formally, let \mathcal{D}' be the subset of nodes in A adjacents to the edges of the cut. For each $d' \in \mathcal{D}'$, d' acts as an R -generalized destination with $out(d') = |\{(d', v), v \in B\}|$. So similarly as for B , A is stable.

3.3.4 $S \cap \mathcal{D}' \neq \emptyset$

If generalized sources s are in \mathcal{D}' , then we can look inside A and consider them as generalized sources with modified arrival rate $in'(s) = in(s) - |\{(s, v), v \in B\}|$, if this is positive. If $in(s) - |\{(s, v), v \in B\}|$ is negative, then we consider s as a generalized destinations. As for Section 3.3.2, we have a feasible arrival rate, and so by induction, the network is stable.

The induction ends when we arrive either into the uniformly unsaturated case that is proved in Section 3.1, or when $A = S$ and $B = \mathcal{D}$. In the latter case, we know that we have stability since the arrival rate is feasible.

4 Conclusion

In this paper, we show that the *LGG* protocol is stable. We conjecture that the protocol stays stable even when an adversary injects some flow that can violate "a little" the feasibility constraint.

Moreover, we conjecture analogous stability results when the network is dynamic. In particular, it seems that our protocol is stable if a flow of value f always exists, at least in the uniformly unsaturated case.

References

- [1] A. Goldberg and R. Tarjan. A new approach to the maximum-flow problem. *Journal of the Association for Computing Machinery*, 35(4):921–940, Oct. 1988.
- [2] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12):1936–1948, Dec. 1992.
- [3] P. Tsaparas. Stability in adversarial queueing theory. Technical report, Master Thesis, Department of Computer Science, University of Toronto, 1997.
- [4] X. Wu and R. Srikant. Regulated maximal matching: A distributed scheduling algorithm for multi-hop wireless networks with node-exclusive spectrum sharing. In *44th IEEE Conference on Decision and Control, and the European Control Conference (CDC-ECC)*, pages 5342–5347, Dec. 2005.



Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399