



Substitution Networks: Performance Collapse due to Overhead in Communication Times

Thiago Abreu, Nghi Nguyen, Thomas Begin, Isabelle Guérin Lassous, Bruno Baynat

► To cite this version:

Thiago Abreu, Nghi Nguyen, Thomas Begin, Isabelle Guérin Lassous, Bruno Baynat. Substitution Networks: Performance Collapse due to Overhead in Communication Times. 4th International Conference on Ad Hoc Networks (ADHOCNETS 2012), Oct 2012, Paris, France. Springer, 111, pp.1-16, 2012, Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering. <10.1007/978-3-642-36958-2_1>. <hal-00763225>

HAL Id: hal-00763225

<https://hal.inria.fr/hal-00763225>

Submitted on 10 Dec 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Substitution Networks: Performance Collapse due to Overhead in Communication Times

Thiago Abreu¹, Nghi Nguyen², Thomas Begin¹, Isabelle Guérin Lassous¹, and Bruno Baynat²

¹ Université Lyon 1 - LIP (UMR ENS Lyon, CNRS, INRIA, UCBL), Lyon, France, {thiago.wanderley, thomas.begin, isabelle.guerin-lassous}@ens-lyon.fr

² Université Pierre et Marie Curie - LIP6, Paris, France {nghi.nguyen, bruno.baynat}@lip6.fr

Abstract. A substitution network is a wireless solution whose purpose is to bring back connectivity or to provide additional bandwidth capacity to a network that just suffered a failure or a dramatic surge in its workload. We analyze the performance of the simplest possible multihop topology for a substitution network, i.e., the *multihop chain* subject to traffic transmitted in both directions. Clearly, the potential capacity of a substitution network, whose technology should be embedded in mobile routers, is very likely to be far much smaller than the prior base network. We investigate the actual performance attained by such a substitution network under various conditions of the chain length and the carrier sensing range. Our results show that the capacity, viz. its maximum attainable throughput, reaches a peak at a given workload and then, for larger values of workload, decreases towards an asymptote which value can be drastically lower than the peak value. We give insights into this performance collapse and show the need for a suitable admission control.

Key words: Substitution networks, Multihop chain, Performance collapse, IEEE 802.11

1 Introduction

A *substitution network* stands for a wireless solution whose purpose is to bring back connectivity or to provide additional bandwidth capacity to a network that just suffered a failure or a dramatic surge in its workload (e.g., flash crowd effect). This latter network is called the *base network*, and it can be based on wired or wireless technologies. It is worth pointing out that, unlike other ad hoc and mesh solutions, a substitution network does not seek to provide new services to customers. Its goal is rather to restore and/or maintain at least some services that were available prior to the base network troubles. As a matter of fact, a substitution network is not a stand-alone network.

Two types of nodes are involved in a substitution network: (i) *Bridge routers*, which are basically gateways interconnecting the base network and the substitution network; (ii) *Mobile routers*, which are the core piece of the substitution network. Their positioning should be done so as to give rise to path(s) that will

route the traffic delivered by the base network through the substitution network. Obviously, bridge routers require a wireless interface to connect to the substitution network, and mobile routers require motion capabilities to move towards their expected position. Wifibots¹ or micro-drones, like for instance AR.Drone 2.0², can be used as mobile routers. Last but not the least, an algorithm should decide where mobile routers should move to.

The concept of a substitution network has been initially proposed in [1] and it is also the core focus of the ANR VERSO RESCUE project (ANR-10-VERS-003)³. In this context, the base network is assumed to operate through wired technologies or through a wireless technology requiring large and fixed facilities. So it is very likely that the attained capacity of the substitution network, whose technology should be embedded in mobile routers, is far much smaller. Such a drop in the available bandwidth has clear implications with regard to the control policy to be implemented on the traffic at bridge routers. However, prior to these operations, it is crucial to position the mobile routers in the best possible way and so to investigate the performance behavior of a substitution network. For instance, investigated factors may include its available bandwidth with regard to the deployed topology, transmission power assigned to router antennas, and buffer size on the wireless interfaces.

The scope of this paper is restricted to the performance analysis of the simplest possible multihop topology for a substitution network, namely the *multihop chain*. The case of a multihop network arranged as a chain topology has been extensively studied, especially when IEEE 802.11 technology in DCF mode is used as MAC protocol. Virtually all, if not all, of these works have dealt with the case of a one-way traffic, viz., traffic transmitted from one endpoint (i.e., bridge routers) of the chain to the other endpoint. Obviously, providing the injected traffic is transported using TCP protocol, then a reverse traffic consisting of ACKs packets is departing from the destination node towards the source node, but its workload is very light and tightly correlated to the data traffic. On the other hand, in the case of a substitution network, the transported traffic is expected to be of roughly equal size in both directions. Therefore, the potential emergence of substitution networks argues the need for new studies that handle the case of a chain topology but subject to a two-way traffic workload and, as far as we know, there is no work providing results and insights for this scenario. In addition, a two-way traffic implies more interferences and collisions in the wireless substitution network that, as shown later, will generate unexpected and critical behaviors.

The paper is organized as follows. Section 2 details the used scenario and the corresponding notation. In Section 3, we discuss the experimental results obtained by simulation. Section 4 brings some insight into the performance collapse potentially exhibited by a substitution network. Section 5 concludes this paper.

¹ <http://www.wifibot.com>

² <http://ardrone2.parrot.com/>

³ <http://rescue.lille.inria.fr/>

2 Scenarios

2.1 Description and assumptions

In this paper, we consider the case of a multihop chain scenario. This is the simplest topology for a substitution network that aims at restoring connectivity to a base network that is in trouble. The number of required hops in the chain depends on the distance in between the two endpoints where the connectivity needs to be brought back. Obviously, performance of the network vary deeply with the number of hops, and hence we consider several lengths of chains. In all tested scenarios, nodes are equidistant (any node is separated from its neighbor(s) by 250 meters).

As said previously, traffic is likely to be exchanged in various directions in the base network. Therefore, it is fair to expect that the traffic sent over the substitution network consists of data transmitted in both directions (from one endpoint of the chain up to the other endpoint). We model the traffic of datagrams in each direction by two independent Poisson processes. The resulting traffic submitted to the substitution network (sum of two independent Poisson processes) is thus also a Poisson process. We refer to the workload as the rate of this resulting process, i.e., the rate of datagrams attempting to access the network at any of the two endpoints. The valued of the workload will be denoted by λ . Hence, at each endpoint, the local workload is of $\frac{\lambda}{2}$. In the aim of assessing the capacity (i.e., the maximum attainable throughput) of a network, UDP traffic appears as a fair choice (unlike TCP, there is no control loop). On top of that, UDP sources may be a suited choice to represent an aggregation of dozens or hundreds of uncorrelated flows. More complex traffic patterns are left for future work.

In our scenario, we rely on-the-shelf wireless technologies for the substitution network. More precisely, we assume that wireless communications are based on IEEE 802.11 and use the DCF mode based on a CSMA/CA approach. Since previous studies has shown that the use of RTS/CTS is ineffective in multihop wireless networks [2, 3], we do not include the case of using RTS/CTS in our scenarios.

We also investigate the impact of different carrier sensing ranges on the multihop chain performance. Since the radio environnement may undergo various propagation properties, we consider both the cases of a carrier sensing range greater than the communication range and of equal size. In the last case a node can only sense communications of one-hop neighbors, whereas in the former case it can also sense communication of two-hop neighbors.

Finally, for sake of simplicity we assume that the physical layer is perfect. It means that any frame transmitted on the radio medium is always received successfully if no collision occurs. Of course, this represents a strong assumption that does not match with reality. However the performance analysis under an ideal physical layer assumption is an unavoidable first step in the quest of better understanding the MAC layer behavior in terms of frames collisions, backoff

freezes and buffer overflows. In addition, we believe that our qualitative observations will not drastically change with a more realistic physical layer. This will be verified in future work.

2.2 IEEE 802.11 DCF main principles

In this section, we remind the main principles of the DCF (Distributed Coordination Function) mechanism of IEEE 802.11 [4], since we rely on precise elements to discuss the obtained results in Section 3. A node that wants to transmit a datagram is required to wait for a fixed period DIFS and another period, called *backoff*, randomly chosen in a contention window. During the backoff period, the backoff is decreased by one whenever the medium is sensed as idle, at the end of a slot duration. If the medium is detected as busy, the backoff counter is frozen, and is resumed as soon as the medium returns to the idle state. When the backoff counter reaches zero, the node starts the transmission of the frame corresponding to the datagram. In 802.11 DCF, (unicast) frames are considered to be successfully transmitted if the sender receives an acknowledgement from the receiver after a fixed and short period SIFS and prior to a given timeout. Otherwise, the sender considers that a collision has occurred. The datagram is then retransmitted by using the same process but with a contention window whose size is doubled (up to a given maximum value).

The maximum number of transmissions for a given datagram is limited to a maximal value defined by the standard. Typically, this value is set to 7. Note that the size of the contention window is reset to its initial value whenever a frame transmission succeeds or when a datagram is dropped because it exceeds the maximal number of retransmissions.

2.3 Notation

We now detail the main notation we use in this paper. Let us consider a multihop chain with n hops (meaning $n + 1$ nodes labelled from 0 to n).

We let $t_{a,b}$ denote the mean throughput of the traffic sent from node a to node b , which is computed as the number of successfully transmitted datagrams (in bits) per unit of time. $T_{a,b}$ indicates the mean throughput of the aggregate traffic exchanged between nodes a and b (in both directions). By definition, we have:

$$(1) \quad T_{a,b} = T_{b,a} = t_{a,b} + t_{b,a} \quad \forall a \text{ and } b \in [0, n]$$

Figure 1 illustrates this notation.

Simple and obvious relations can be established between these latter quantities, and come from the fact that traffic arriving to node n (resp. node 0) corresponds to datagrams that have been sent by node 0 (resp. node n) and that have not been lost in between both:

$$(2) \quad T_{0,n} = t_{1,0} + t_{n-1,n}$$

$$(3) \quad t_{a,a+1} \geq t_{a+1,a+2} \quad \forall a \in [0, n - 2]$$




Fig. 1. Notation used in the paper

(4) $t_{0,n} = t_{n-1,n}$ and $t_{n,0} = t_{1,0}$

By the system symmetry, we also have:

(5) $T_{0,1} = T_{n-1,n}$ and $t_{a,a+1} = t_{n-a,n-a-1} \quad \forall a \in [0, n - 1]$

Similarly, we denote by $p_{a,a+1}$ the probability that a frame transmitted from node a to node $a + 1$ is lost because of a collision. $P_{a,a+1}$ then represents the probability that a frame transmitted between nodes a and $a + 1$ (in both directions) gets corrupted due to a collision. Finally, we let $q_{a,a+1}$ denote the probability that a datagram transmitted from node a to node $a + 1$ is lost because it experienced seven consecutive collisions and $Q_{a,a+1}$ the probability that a datagram transmitted between nodes a and $a + 1$ (in both directions) is lost.

The next two following equations are approximately satisfied:

(6) $q_{a,a+1} \simeq (p_{a,a+1})^7 \quad \forall a \in [0, n - 1]$
 (7) $Q_{a,a+1} \simeq (P_{a,a+1})^7 \quad \forall a \in [0, n - 1]$

3 Simulation

We assess the performance of the multihop network through extensive discrete-event simulations using Ns2.35 for various lengths of the chain, viz. $n = 3, 4$ and 5 . We relate the parameters values we use in Table 1.

3.1 Case of a 3-hop chain

A - Carrier sensing range greater than transmission range

We start our analysis with the case of $n = 3$ hops, and a carrier sensing range greater than transmission range. Let us recall that this only means that a node can sense communication of two-hop neighbors, e.g., node 1 can sense communications of nodes 0, 2 and 3, but node 0 can only sense communications of nodes 1 and 2. Figure 2 shows the related results as a function of λ . Remind that a workload of $\lambda/2$ is injected both on nodes 0 and n .


Figure 2(a) indicates the values of $t_{a,a+1}$ for each link, i.e., the mean traffic throughput from node a to node $a + 1$ (viz. in one direction) and $t_{0,n}$, i.e., the

Propagation model	2-Ray Ground
Physical rate	11 Mb/s
Transmission range	399 m
Carrier sensing range	709 m or 399 m
Distance between neighbor nodes	250 m
RTS/CTS	Disabled
Traffic	Two Poisson sources with UDP. One for each direction
Workload λ	From 0.1 to 10 Mb/s with steps of 0.1 Mb/s
Packet size	1500 bytes
SIFS	10 μ s
DIFS	50 μ s
Backoff time slot	20 μ s
Retransmission limit	7
Contention window size (min, max)	32, 1024
Node buffer size	50 packets
Simulation time	150 s


Table 1. Experimental setup for the simulations

mean traffic throughput from node 0 to node n . First, as λ increases from 0.1 up to around 2Mb/s, so do $t_{a,a+1}$ and $t_{0,n}$. Then, the values of $t_{a,a+1}$ and $t_{0,n}$ tend to be flattened at a value often close (sometimes less) to their maximum value. Note that, in this area, we have: $t_{a,a+1} \geq t_{a+1,a+2} \forall a \in [0, n - 2]$, which is in line with relation (3) from Section 2.3. This behavior ensues from the combined effect of collision occurrences and buffer overflows. As shown in Figure 2(b), $p_{0,1}$, i.e., the probability that a frame transmitted from node 0 to node 1 experiences a collision, is far much larger (around ten times) than $p_{1,2}$ and $p_{2,3}$ because of node 3 that acts as a hidden terminal. However, the values of $q_{0,1}$ and a fortiori those of $q_{1,2}$ and $q_{2,3}$, are almost equal to 0 (consistent with relation (6)), meaning that virtually no datagram will be lost due to consecutive frame collisions. Yet this significant value of $p_{0,1}$ tends to “elongate” the time the node 0 needs to successfully transmit a datagram to node 1 but also the time that nodes 1 and 2 take to forward the datagram due to backoff freezes. Then, given the finite size of buffers at each node, this elongation in the time needed for serving a datagram may ultimately cause datagram losses by overflow. This is in line with the decreasing values of $t_{a,a+1}$ in Figure 2(a). Finally and unsurprisingly, the values of $t_{0,n}$ match those of $t_{2,3}$ here, which complies with the relation (4).


We turn now to Figure 2(d). Let us remind that $T_{a,a+1}$ denote the mean throughput of the aggregate traffic exchanged between nodes a and $a + 1$ (in both directions). First, $T_{0,1}$ and $T_{2,3}$ are identical due to the intrinsic system symmetry (see relation (5)). Second, as expected, $T_{1,2}$ is slightly less than $T_{0,1}$ since datagram losses occur due to buffer overflow at nodes 1 and 2. By the way, this figure is also in line with relation (1). For instance, $t_{1,2}$ (roughly equal to 1.2 Mb/s from Figure 2(a)) summed up to $t_{2,1}$ (equal to $t_{1,2}$, see relation (5)) comes equal to $T_{1,2}$. Third, $T_{0,3}$ is steadily lower than $T_{0,1}$, $T_{1,2}$ and $T_{2,3}$, which




(a) Throughput on each link on ascending nodes direction



(b) Frame collision probability on each link on ascending nodes direction



(c) Lost datagrams by buffer overflow at each node and for the overall network



(d) Aggregate throughput on both direction for each link and for the overall network


Fig. 2. $n = 3$ hops and carrier sensing range greater than transmission range

agrees to the fact that $T_{0,3}$ is made up as the sum of the two smallest $t_{a,a+1}$ from Figure 2(a), i.e., $t_{2,3}$ and $t_{1,0}$. Last but not least, it is worth pointing out that the evolution of $T_{0,3}$ hits its top value around $\lambda = 2.1$ Mb/s before decaying slightly. As shown below this unexpected behavior will be accentuated in other examples.


B - Carrier sensing range equal to transmission range

We now turn to the case where the carrier sensing range equals the transmission range, i.e., a node cannot sense communication of more than one-hop neighbors. n is kept to 3 hops. We report in Figure 3 the related results (again as a function of λ).


Figure 3(a) shows the obtained results for $T_{0,1}$, $T_{1,2}$, $T_{2,3}$ and $T_{0,3}$. Same remarks can be made here as in the previous case (with a carrier sensing range greater than transmission range). Yet the main difference is quantitative. Obtained values tend to be smaller than previously reported values. The aggregate throughput attained by the network converges towards 1.2 Mb/s with a peak



(a) Aggregate throughput on both direction for each link and for the overall network



(b) Frame collision probability for each link on ascending nodes direction



(c) Lost datagrams by buffer overflow at each node and for the overall network

Fig. 3. $n = 3$ hops and carrier sensing range equal to transmission range

close to 1.6 Mb/s. First, it is worthwhile pointing out that the attained throughput for high traffic is almost half of the one obtained in the previous example (1.2 Mb/s instead of 2 Mb/s). Second, the difference between the asymptote and the peak is also bigger than in the previous example and becomes now significant. It is a consequence of a smaller carrier sensing range that implies more collisions when traffic is high. This behavior may indeed be tied to the rise of the frame collision probabilities $p_{0,1}$ and $p_{1,2}$, which are now laying at much higher values (see Figure 3(b)), since two nodes, viz. nodes 0 and 1, instead of one, are actually subject to hidden terminals. Then, not surprisingly, as frame collision probabilities increase, so do the time needed by a node for serving a datagram which may ultimately cause datagram losses by buffer overflows. Datagram losses are reported in Figure 3(c).

It thus appears that, at least in this case, it could be worth controlling the workload rates at the input of the network so as to maximize the overall throughput performance of the network.

3.2 Case of a 4-hop chain

A - Carrier sensing range greater than transmission range

We carry on our analysis with the case of $n = 4$ hops, and a carrier sensing range greater than transmission range. Associated results are reported in Figure 4.


Figure 4(a) relates the values of $T_{0,1}$, $T_{1,2}$, $T_{2,3}$, $T_{3,4}$ and $T_{0,4}$. These results somehow differ from those obtained for $n = 3$ (see Figure 2(d)). On the one hand, $T_{0,1}$ and $T_{3,4}$ keep increasing towards an asymptote close to 4 Mb/s. On the other hand, $T_{1,2}$ and $T_{2,3}$ exhibit a peak value around $\lambda = 2.2$ Mb/s and then steadily decrease. Then, not surprisingly, $T_{0,4}$, which represents the aggregate throughput of the network, reaches its peak around $\lambda = 1.9$ Mb/s before decreasing significantly towards 0.5 Mb/s. It is worth noticing that in this case the ratio between the maximum value of $T_{0,4}$ and its asymptotic value is close to 3. (This ratio was only around 1.1 for $n = 3$). In this example, the need for an admission control in order to avoid this huge drop in throughput appears to be crucial.

To go further, we look at the frame collision probability in Figure 4(b). Unlike the case of $n = 3$ (see Figure 2(b)), $p_{0,1}$ exhibits a peak value here around 1.8 Mb/s and then decreases gradually. In addition, again as opposed to the case of $n = 3$, there is another collision probability that is non-null, viz. $p_{1,2}$. This growth of $p_{1,2}$ is mostly due to the presence of a hidden terminal for node 2 when $n = 4$.


The collision probabilities ensue losses of frames, but with such levels of probability, it is very unlikely that a datagram experiences seven consecutive collisions and then gets lost. Datagrams may rather be lost when attempting to access the buffer of nodes when λ is large due to buffer overflows. Figure 4(c) reports these datagram losses at each node as a function of λ . First, as opposed to $n = 3$, losses are now mainly occurring on intermediate nodes, viz. nodes 1 and 3 (though nodes 0 and 4 are still experiencing losses but at a lesser degree). Second, as expected, the total amount of losses in the network, designated by the red curve, exhibits a first regime with no losses going up to around 2 Mb/s, and a second regime with increasing losses. Perhaps more interestingly, on this latter regime, the slope of the curve is greater than 1 and close to 2 when the traffic is slightly higher than 2 Mb/s. This means that, just beyond 2 Mb/s, as the workload of the network is increased by A , then the total amount of losses is roughly increased by $2 \times A$. This somehow counterproductive phenomenon leads to the decay exhibited by $T_{0,4}$ on Figure 4(a).

B - Carrier sensing range equal to transmission range


Figure 5 reports the results obtained for $n = 4$ but with a carrier sensing range equal to transmission range. Because of the network symmetry, we get $T_{0,1} = T_{3,4}$ and similarly, $T_{1,2} = T_{2,3}$ as shown by Figure 5(a). This latter figure also shows that $T_{0,4}$, i.e., the aggregate throughput of the network, hits its maximum value around 1.2 Mb/s and then decreases to its asymptotic value. It is worth noting that having carrier sensing range equal to transmission range yields a smaller maximum value of $T_{0,4}$ as compared to having a carrier sensing



(a) Aggregate throughput on both direction for each link and for the overall network.



(b) Frame collision probability for each link on ascending nodes direction.




(c) Lost datagrams by buffer overflow at each node and for the overall network.


Fig. 4. $n = 4$ hops and carrier sensing range greater than transmission range

range greater than transmission range (viz. 0.3 Mb/s), yet the asymptotic value of $T_{0,4}$ for large value of workload λ is slightly higher (viz. 0.2 Mb/s). However the ratio between the peak and the asymptote remains greater than 2 and thus controlling the workload rate at the input of the network would still be very productive.

Here also, datagram losses occur overwhelmingly due to buffer overflows (and not because of seven consecutive collisions). Figure 5(b) relates the amount of losses at each node. As opposed to previous cases, losses are somehow shared among the nodes (with the exception of node 2). Finally, the total amount of losses in the network, represented by the red curve on Figure 5(b), shows that the network experiences virtually no losses up to a workload $\lambda = 1.5$ Mb/s. Afterwards, the curves starts increasing with a slope starting to 1.5 (for $\lambda < 3$ Mb/s) and decaying to 1. This behavior perfectly complies with Figure 5(a) where one can observe that for increasing values of λ within $[1.5, 3]$ Mb/s, the aggregate throughput of the network, i.e., $T_{0,4}$, tends to decrease. On the other hand, for λ values greater than 3 Mb/s, $T_{0,4}$ holds constant.



(a) Aggregate throughput on both direction for each link and for the overall network



(b) Lost datagrams by buffer overflow at each node and for the overall network

Fig. 5. $n = 4$ hops and carrier sensing range equal to transmission range




Fig. 6. $n = 5$ hops and carrier sensing range greater than transmission range




Fig. 7. $n = 5$ hops and carrier sensing range equal to transmission range

3.3 Case of a 5-hop chain

A - Carrier sensing range greater than transmission range

Finally we consider the case of $n = 5$ hops. In Figure 6, we represent the aggregate throughput for each link when carrier sensing range is greater than transmission range. As shown on this figure, $T_{2,3}$ plays as the bottleneck link for this network. Note that the aggregate throughput of the overall network, i.e., $T_{0,5}$, tends to stand below $T_{2,3}$ since datagram losses may still occur after having passed the bottleneck link. This figure also clearly states that $T_{0,5}$ hits its maximum value for a moderate value of λ (around 1.5 Mb/s). Afterwards, $T_{0,5}$ decays gradually towards a drastically small value less than 0.1 Mb/s. Here, more than never, an admission control would be of high productivity.

B - Carrier sensing range equal to transmission range

Figure 7 reports the simulation results in the case where carrier sensing range and transmission range match. First, all $T_{a,a+1}$ and $T_{0,5}$ exhibit an asymptotic

value for large values of λ . $T_{0,5}$, which represents the aggregate throughput of the overall network, peaks at 1.2 Mb/s for $\lambda = 1.2$ Mb/s, and then decreases gradually towards its asymptotic value of 0.4 Mb/s.

4 Insights into results

A key result of Section 3 is that for all tested values of n the aggregate throughput of the network, i.e., $T_{0,n}$, always peaks at a given workload value λ . For smaller values of λ , and more surprisingly for greater values of λ , $T_{0,n}$ reaches lower values. Such a performance collapse clearly argues for controlling the rate at which the workload is introduced to the network. Yet before investigating this in more detail, we focus the paper on the reasons that cause this peak to occur.

Overall, the network can be viewed as a black box system that processes datagrams requesting service. Its input corresponds to the datagrams arriving rate, viz. the workload λ , and its output represents the rate of successfully transmitted datagrams through the chain, viz. the aggregate throughput of the network $T_{0,n}$. This is illustrated by Figure 8.




Fig. 8. Network as a black box system

We also represent potential losses on this figure since the network can experience datagram losses if the buffer of a node is found full upon a datagram arrival. Remind that our experiments demonstrated that virtually no datagrams were lost because of 7 consecutive frames collisions.

We denote by t_S the time needed by the network to successfully process a datagram, which corresponds to the sojourn time spent by a datagram within the network. This service time includes potential waiting times at buffers, backoff overheads, frames transmissions and retransmissions. Let us emphasize that t_S is not a constant time here, but a function of λ . As λ increases, the waiting time spent by datagrams queueing in buffer nodes tends to increase as well as the backoff overheads, causing ultimately a growth of t_S . The system is said to




Fig. 9. Evolution of the service time t_S against the workload λ for $n = 5$

undergo an elongation of its service time. Such a feature is quite common for systems where requests need to access a shared resource to be served. To better understand its evolution, we capture its values in our different simulations. Figure 9 represents the values of t_S against λ for all previous cases (n going from 3 to 5 hops). For $n = 5$ and a carrier sensing range greater than (respectively,

equal to) the transmission range, we observe that t_S is close to 0.001 s for low values of λ and then increases steadily, reaching values around 5 s (respectively, 2 s). Similar remarks hold for other values of n .

The capacity of a system usually refers to its maximum attainable throughput. In our case, for a wireless network with n hops, the system capacity varies with λ . This change can be observed in Figures 2(d), 3(a), 4(a), 5(a), 6 and 7 where the network capacity tends to decrease significantly once the peak value has been reached. In a broader way, the system capacity highly depends on its service time but also on other factors. Another key factor is the DOP (degree of parallelism). The parallel processing stems from potential pipelining in the datagram service. For instance, as node a handles its datagram (viz. transmits the corresponding frame on the channel), node $a + 4$ can also be in transmission. On top of this, another well-known pipelining effect may occur when two neighbor nodes decrement simultaneously their backoff. We denote by D this degree of parallelism. Typically, we have $D \geq 1$ but it is unclear how to get its exact value. Based on simple queueing theory elements, it follows that the system capacity is equal to $\frac{D}{t_S}$.

The aggregate throughput of the network, i.e., $T_{0,n}$, can obviously not exceed $\frac{D}{t_S}$, neither λ (it can be less than λ due to datagram losses). Hence we have:

$$T_{0,n} \leq \min\left(\lambda, \frac{D}{t_S}\right)$$




Fig. 10. Schematic view of $T_{0,n}$ against λ

Note that when $\lambda > \frac{D}{t_S}$, then the datagram loss rate increases dramatically. We schematically plot in Figure 10 the evolution of the upper bound of $T_{0,n}$ assuming for t_S an evolution similar to the ones exhibited on Figure 9.

The obtained curve for the upper bound of $T_{0,n}$ exhibits a peak value corresponding to the intersection point of the function λ and the one designating the system capacity. This is in line with our simulated experiments (see Figures 2(d), 3(a), 4(a), 5(a), 6 and 7).

5 State-of-the-Art

In [5, 6], the authors study, by simulation, the performance of IEEE 802.11 DCF in multihop chains when the injected traffic consists of TCP flows. They point out the TCP instability and the fairness issues that may arise with this topology. Their simulations always include the use of RTS/CTS, and the TCP flows, which compete to access the network resource, may only go through the network on a segment of the chain (whereas in the case of a substitution networks, it is very likely that the traffic is exchanged from one endpoint to the other). During the same year, the authors of [7] study the capacity of multihop wireless networks

based on IEEE 802.11 DCF and, specifically, the capacity of multihop chains. The performance are studied using the simulator Ns2. The authors turn on the RTS/CTS mechanism, set the physical rate of wireless links to 2 Mb/s, and assume that a single flow is sent through this topology (in fact the flow profile is not specified in the article). The simulation results show that, when the flow rate becomes too large, the throughput achieved along a multihop chain tends to decrease with the number of hops of the chain and seems to converge towards a bound. They also show that the capacity of the chain can be attained for a specific flow sending rate and that the the throughput obtained on the chain slightly decreases and becomes smaller than the capacity when the flow sending rate is increased. Other topologies are tested but none of them corresponds to our scenario.

In [8], the authors carry out real experiments on a multihop chain and comment the obtained performance. The chain consists of at most 4 hops and one CBR/UDP flow is transmitted from one endpoint of the chain to the other endpoint at a saturating rate (i.e., there is always a packet to send at the source), with a physical layer of 2 Mb/s and without the use of RTS/CTS. The results show that the obtained throughput on the chain is very low (and smaller than the one obtained by simulation) and very unstable. They also show that the third and fourth links have lower performance than the first and second links. These experiments only consider one saturating one-way flow, which is different from the scenario we consider.

More recently, the multihop chain topology has been investigated in details. In [9], the authors provide a throughput analysis of multihop chains with one transmitted flow. The analysis is carried out both by simulation and analytically and shows that the load injected into the chain must be controlled in order to reach the optimal overall throughput. This analysis focuses on interactions between hidden nodes. A similar study is provided in [10]. In [11], the authors study the performance of a multihop chain under different physical rates, with and without the use of RTS/CTS, and with one-way flow. The results show that using high physical rates do not necessarily lead to the best performance in terms of end-to-end throughput and end-to-end delay, particularly when RTS/CTS are used.

In [12], the authors study potential interactions that can arise between links of a chain. They evaluate, by simulation and with real experiments, the impact of the different possible interactions on the overall chain throughput. The obtained results show that some kinds of interactions give rise to better performance and that the kind of interaction at the beginning of a chain has more impact on the performance than at the end of the chain. The results also show that, beyond a given sending rate of the injected flow, the overall throughput in the chain slightly decreases. This work had been started in [13] with a less realistic model for the packet reception. This study do not consider the targeted scenario of this article. In these works, CBR/UDP flows are transmitted. These studies are extended to TCP flows in [14]. But as explained in Introduction, even though a TCP flow implies two flows at the transport level, one in each direction of

the multihop chain, these two flows have dependencies and asymmetric sending rates, as opposed to our work.

6 Conclusion

In this paper, we analyze the performance of the simplest possible multihop topology for a substitution network, i.e., the *multihop chain* subject to traffic transmitted in both directions. Clearly, the potential capacity of a substitution network, whose technology should be embedded in mobile routers, is very likely to be far much smaller than the prior base network. We investigate the actual performance attained by such a substitution network under various conditions of the chain length and the carrier sensing range. Our results confirm that the potential capacity of a substitution network, viz. its maximum attainable throughput, tends to be low. More interestingly, we show that its capacity exhibits a peak value around a given workload rate and then, for larger values of workload, decreases gradually. We provide insights into this performance collapse. Future work will be devoted to automatically detect this optimal running point of the network, and so control the rate at which the workload is introduced to the network.

Acknowledgments

This work was partially funded by the French National Research Agency (ANR) under the project ANR VERSO RESCUE (ANR-10-VERS-003).

References

1. Razafindralambo, T., Begin, T., Dias De Amorim, M., Guérin Lassous, I. Mitton, N., Simplot-Ryl, D.: Promoting Quality of Service in Substitution Networks with Controlled Mobility. AdHocNow 2011, Germany (2011)
2. Xu, K., Gerla, M., Bae, S.: Effectiveness of RTS/CTS handshake in IEEE 802.11 based ad hoc networks. Ad Hoc Networks, Issue 1, pages 107-123, Elsevier (2003).
3. Ray, S., Starobinski, D.: On False Blocking in RTS/CTS-based Multi-hop Wireless Networks. IEEE Transactions on Vehicular Technology, Vol. 57, Number 2, March 2007.
4. L.S.Committee, IEEE Computer Society: ANSI/IEEE Std 802.11: Wireless LAN Medium Access Control(MAC) and Physical Layer (PHY) Specifications, 2007.
5. Xu, S., Saadawi, T.: Does the IEEE 802.11 MAC protocol work well in multihop wireless ad hoc networks? IEEE Communications Magazine, Vol. 39, No. 6., pp. 130-137 (2001).
6. Xu, S., Saadawi, T.: Revealing the problems with 802.11 medium access control protocol in multi-hop wireless ad hoc networks. Computer Networks, Volume 38 Issue 4, pp 531-548, Elsevier (2002).

7. Li, J., Blake, C., De Couto, D. S. J., Lee, H. I., Morris, R.: Capacity of Ad Hoc Wireless Networks. MOBICOM, pp. 61-69 (2001).
8. Dhoutaut, D., Guérin Lassous, I.: Performance of a multi-hops configuration with 802.11: from simulation to experimentation. PIMRC, Barcelona, Spain, (2004).
9. Ng, P. C., Liew, S. C.: Throughput Analysis of IEEE 802.11 Multi-Hop Ad Hoc Networks. IEEE Transactions on Networking, Vol 15 , Issue 2, pp 309-322 (2007).
10. Yoo, J.-Y., Kim, JW.: Maximum End-to-End Throughput of Chain-Topology Wireless Multi-Hop Networks. Wireless Communications and Networking Conference (WCNC), pp 4279-4283 (2007).
11. Li, F. Y. L., Hafslund, A., Hauge, M., Engelstadt, P., Kure, O., Spilling, P.: Does Higher Datarate Perform Better in IEEE 802.11-based Multihop Ad Hoc Networks? Journal of Communications and Networks, Vol. 9, Number 3, pp 282-295 (2007).
12. Razak, S., Kolar, V., Abu-Ghazaleh, N. B., Harras, K. A.: How do wireless chains behave?: the impact of MAC interactions. MSWiM, pp 212-220 (2009).
13. Razak, S., Abu-Ghazaleh, N. B.: Self-interference in Multi-hop Wireless Chains: Geometric Analysis and Performance Study. ADHOC-NOW, pp 58-71 (2008).
14. Majeed, A., Abu-Ghazaleh, N. B., Razak, S., Harras, K. A.: Analysis of TCP performance on multi-hop wireless networks: A cross layer approach. Ad Hoc Networks, Volume 10, Issue 3, pp 586-603, Elsevier (2012).