A Study of the Performance of Cooperative Caching in Static Ad Hoc Networks

Francisco J. González-Cañete and Eduardo Casilari
Department of Electronic Technology
University of Málaga
Málaga, Spain
{fgc, ecasilari}@uma.es

***Abstract*—In this paper, we evaluate the performance of the CLIR (Cross-Layer Interception and Redirection) cooperative caching scheme for ad hoc networks in static grid ad hoc networks. This caching scheme implements a local cache in every node of the network allowing the nodes to work as request interceptors for the rest of the nodes. In addition, it also implements a redirection cache that stores information about the location of the documents in the network in order to redirect the requests to nodes that are situated closer than the servers. Finally, a piggy-backing technique is incorporated to the routing protocol with the aim of finding the documents in the network while the routes to the servers are created. By means of simulations, we evaluate the mean traffic generated in the wireless network, the delay perceived by the users, the percentage of failed searches, the mean number of retrieved documents, the local and remote cache hits and the mean percentage of redirection hits as a function of the mean time between requests, the Time To Live (TTL) of the documents, the traffic pattern and the cache sizes. We compare the performance of our proposal with another five cooperative caching schemes as well as the option of no using a caching scheme. The simulation results show that our proposal outperforms the other caching schemes in terms of the studied parameters. In addition, we compare the redirection caching scheme of our proposal to the redirection policies implemented by other caching schemes.**

Keywords-cooperative caching; grid; ad hoc network; redirection cache.

I. INTRODUCTION

The CLIR (Cross-Layer Interception and Redirection) caching scheme was evaluated in static grid Ad Hoc networks in [1]. Therefore, the aim of this work is to extend this performance evaluation by analysing new metrics: the mean number of retrieved documents, the percentage of cache hits (local and remote). This will offer a more concise comparison of the performance of CLIR to other caching schemes. Additionally, the redirection technique will be also evaluated as a metric specifically developed to measure its performance called redirection cache hit.

The aim of a caching scheme is to reduce the traffic generated in the network, as well as the delay perceived by the users and the servers' load [2]. The reduction of the traffic in a wireless network also decreases the probability of collisions and interferences, and hence, the probability of packet loss. Reducing the delay perceived by the users when they request documents improves the user experience and makes the network more attractive to be used. Finally, as a consequence of the caching mechanism, the document requests can be served by other nodes in the wireless network instead of the servers. In a very loaded network, the servers could be a bottleneck as all the requests are sent to them. The caching mechanism mitigates this effect by moderating the overload of the servers so they can reply more requests.

MANETs (Mobile Ad Hoc NETworks) [3] were proposed as a solution for deploying communication applications in places where a wired network was not available. Unfortunately, they have some limitations:

- Restricted hardware capabilities. Some light weight devices are constrained in their processing and computing capabilities.
- Limited batteries. Mobile devices operate with batteries. In order to maximize their lifetimes, the number of messages that they generate should be moderated.
- Scarce bandwidth. Wireless medium has restricted bandwidth so signaling traffic should be minimized.
- Temporary connection to external networks. The integration of MANET into external networks is guaranteed through Gateways. However, the mobility of the MANET may provoke the Gateway to be temporarily unavailable.

Although many cooperative caching schemes have been proposed for MANETs, they have not been evaluated for static ad hoc network, that is, wireless networks where the nodes do not move. The objective of this work is to evaluate the performance of different caching schemes proposed for MANETs in static grid network scenarios.

The rest of this document is organized as follows. In Section II, the related work about cooperative caching schemes for MANETs is presented. In Section III, the proposed caching scheme is described. Section IV defines the system model and shows the performance evaluation of the caching schemes. Finally, Section IV enumerates the main conclusions of this work.

II. RELATED WORK

The cooperative caching schemes for ad hoc networks can be classified into four groups: broadcast-based, information-based, role-based and direct-request. The broadcast-based caching schemes employ broadcast messages as the first choice in order to find the documents in the network. These broadcast messages can be sent to the entire network, as in the case of MobEye [4]. Other schemes such as SimpleSearch [5], follow a more restrictive approach that limits the distance of the messages to four hops. ModifiedSS [6] is an evolution of SimpleSearch that employs GPS (Global Positioning System) in order to send the requests to the direction where the servers are located. Similarly, the caching scheme proposed by Moriya in [7] sends the broadcast messages to the neighbourhood so that, if the document is not found, the request is transmitted to the server.

The information-based cooperative caching schemes employ information of the location of the documents in the network. Nodes obtain this information by analysing the messages that they forward. As examples of this category of caching schemes we can mention: DGA (Distributed Greedy Algorithm) [8], Wang [9], Cho [10] and POACH (POware Aware Caching Heuristic) [11].

Under a role-based caching scheme, each node in the wireless network has a predefined role. That is, they can be caching nodes, requesting nodes, coordinator nodes, gateway nodes, etc. The role-based caching schemes are usually applied to cluster networks. CC (Cluster Cooperative) [12] and Denko [13] are examples of this kind of caching policy.

Finally, the direct-request caching schemes directly send the requests to the server with the hope of being served by an intermediate node in the route from the requester to the server. The proposal by Gianuzzy in [14] is an example of this kind of caching schemes.

However, the groups in this classification of caching schemes are not mutually exclusive. Thus, the caching schemes COOP [15], ORION (Optimized Routing Independent Overlay Network) [16], IXP/DPIP (IndeX Push/Data Pull/Index Push) [17] and COCA (COoperative CAching) [18] are schemes that employ network information and broadcast requests. On the other hand, COACS (Cooperative and Adaptive Caching System) [19] and GROCOCA (GROup-based COoperative CAching) [20] are role-based caching schemes that also utilize information obtained from the network. In addition, CacheData, CachePath, HybridCache [21] and GroupCaching [22] are direct-request caching schemes that also employ the location information. Finally, ZC (Zone Cooperative caching) [23] and Sailhan [24] use direct requests and broadcast requests depending on some heuristic.

The CLIR cooperative caching scheme was proposed in [25]. It can be classified as a direct-request and information-based cooperative caching scheme. The main novelty of CLIR is the implementation of a cross-layer interception cache technique as well as the optimization of the redirection technique. Its performance was evaluated for MANETs and compared to other five cooperative caching schemes. The objective of this paper is to study the performance of CLIR in a static grid ad hoc network and compare this performance with other caching schemes.

Figure 1 summarizes the classification presented in this section.

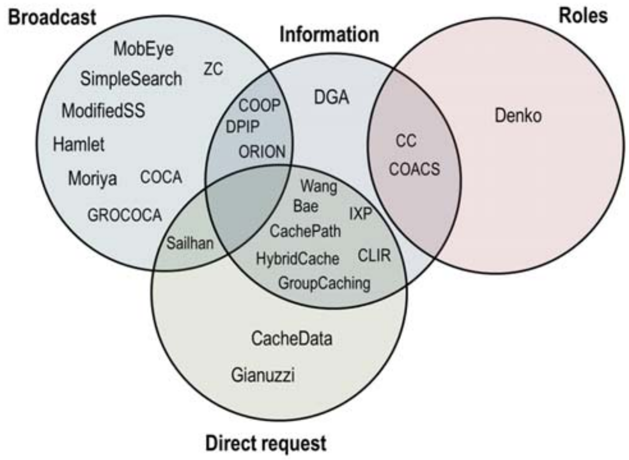

Figure 1. Cooperative caching schemes classification.

III. CACHING SCHEME

The proposed caching scheme works using a request/reply protocol very similar to HTTP (Hyper Text Transfer Protocol) as utilized in the caching schemes proposed in Section II. Consequently, the nodes request documents (information, data, etc.) to the data server. The data servers, as the HTTP servers, reply with a message containing the document requested. The data servers can be other nodes in the wireless ad hoc network or even external servers that are accessed through a Gateway.

CLIR implements a local cache in every node in the network. This local cache is managed using the LRU (Least Recently Used) replacement policy. Using this cache, every node stores the received documents. Therefore, further requests to the same document will be resolved by the local cache. This is called a local cache hit. As the requests must be forwarded hop by hop from the requester node to the server node, the intermediate nodes in the route from the source to the destination of the requests can reply directly if the requested document is stored in their local cache. This is called an interception cache hit.

When the route from the source node of the request to the destination node has not been created, CLIR utilizes the routing protocol to piggy-back the request in the routing protocol messages. By using this technique, the routing protocol is able to create the route to the destination node and search for the requested document at the same time. If any node that receives the route request message has a copy of the requested document in its local cache, it will reply using the route reply message informing that this node has a copy of the document. When the requesting node receives the route reply message, the route between both nodes has just been created so the requester will forward the request to the node that has the copy of the document. This is called a cross-layer interception hit. This mechanism allows finding the documents in the network even if the servers are temporarily unavailable. On the other hand, it also finds the documents in nodes located closer than the servers, reducing the delay and the network load as the number of messages sent are reduced. This kind of technique requires the implementation of an on demand routing protocol as the piggybacking messages are sent on every document request if a route to the destination is not available at this moment.

CLIR also implements a redirection cache that stores information about where the documents are located in the network. In order to obtain this information, the nodes analyse the request and reply messages they forward. The redirection cache manages information about the source of the requests and the corresponding replies. It also stores the number of hops and the TTL of the documents. This TTL is employed to set the validity of the information stored in the redirection cache as the documents will be obsolete after this time. Additionally, the redirection cache also takes into consideration the mean time the documents are stored in the caches in order to avoid the redirection of a request to a node that has evicted the document from its local cache. Consequently, the redirection cache estimates the time that the documents are stored in the local caches calculating the mean time the documents are stored in its own local cache. Moreover, the expiration time assigned to the information stored in the redirection cache is the minimum between this estimated time and the TTL of the document.

The redirection cache is managed by means of two LRU lists, one for the information of those documents whose TTL is known (called *KNOWN_TTL_LIST*) and the other with the documents with an unknown TTL (called *UNKNOWN_TTL_LIST*). The TTL of a document is unknown until the reply is forwarded by the node. The memory for each structure is dynamically assigned although the memory space reserved for both structures is set to a constant. The information obtained from the messages adds or updates the data to the lists. If an entry of the *UNKNOWN_TTL_LIST* structure is updated with the TTL of a document, it will be transferred to the head of the *KNOWN_TTL_LIST*. When a new entry must be stored and there is not enough room because the reserved storage space is full, the oldest entry in the *UNKNOWN_TTL_LIST* structure is evicted. If this list is empty, the oldest entry in *KNOWN_LIST_LIST* will be deleted. An entry is also deleted if the information is obsolete because all the associated TTLs have expired.

When a node receives a request and the redirection cache contains information of a node that is closer to the original destination of the request, the request is forwarded to this closer node. When the node to which the request has been redirected receives the message, it replies with the document. This is called a redirection cache hit. In the case that the redirected node has evicted the document from its local cache, a redirection error message is sent to the redirection node in order to update the information of the redirection cache.

Finally, CLIR also implements the storage of the replied document in the node located in the middle of the route from the source and destination of the reply proposed in [9]. So, the documents can be easily disseminated along the network. In order to avoid the excessive replication of documents, this mechanism is performed if the distance between both nodes is greater than four hops.

For more details on the implementation of CLIR refers to [25].

IV. PERFORMANCE EVALUATION

In order to evaluate the performance of the proposed cooperative caching scheme we have implemented CLIR using the NS-2.33 [26] network simulator. Additionally, for comparison purposes, the cooperative caching schemes MobEye, HybridCache, COOP, DPIP and SimpleSearch have also been implemented. The no cache (NC) option has also been taken into consideration, that is, the case where no caching scheme is considered. Every point represented in the figures shown in this paper corresponds to the mean obtained value of five simulations using the same parameters but changing the seed. Depending on the simulation, the analysed variable is changed while the rest of the parameters are set to a default value. All figures include a confidence interval of 95% for each performance parameter.

A. Simulation model

Table I summarizes the main simulation parameters. We suppose that the nodes in the ad hoc network do not move. Depending on the evaluated configuration, nodes form a regular grid of 5x5, 7x7 or 9x9 nodes. Moreover, the nodes located in the corners of the simulation area, that is, in the positions (x,y)=(0,0) and (x,y)=(1000,1000), are considered to behave as Data Servers (*DS*). There are 1000 different documents stored in the *DSs*. For simulation simplicity, we have considered a numeric identification for each document although the caching scheme can be extended to manage complete URLs. In order to distribute the traffic along the network, the documents with even identification are located in one server while the documents with odd identification are stored in the other *DS*. In addition, every document has an associated TTL (Time To Live) that determines when it expires, and hence, it is considered obsolete. The expired documents stored in the local caches are deleted in order to free storage space. The TTL of the documents follows an exponential distribution with a mean time between 250 and 2000 seconds. In this way, low and high TTL variability is modeled. Moreover, the infinite TTL is also analyzed, that is, the case where the documents never expire. Additionally, we consider the size of the documents to be constant and equal to 1000 bytes.

Every node that is not a server is programmed to generate requests to the servers during the simulation time. When a request is served, another request is generated after a waiting time period. If the request is not served after a predefined timeout, the request is sent again. The waiting time between requests follows an exponential distribution with mean values between 5 and 50 seconds (the default value is 25 seconds). If the waiting time is modified, a wide range of nodes activity can be evaluated and, consequently, the traffic load of the network. The documents that are not served before the timeout is triggered are requested again. This timeout is defined to be constant with a value of 3 seconds.

TABLE I. SIMULATION PARAMETERS

Parameter	Default values	Other utilized values
Simulation area (square meters)	1000x1000	
Number of nodes	49	25-49-81
Number of Servers	2	
Number of documents	1000	
Document size (bytes)	1000	
Timeout (s)	3	
TTL (s)	2000	250-500-1000-2000-∞
Mean time between requests (s)	25	5-10-25-50
Traffic pattern (Zipf slope)	0.8	0.4-0.6-0.8-1.0
Replacement policy	LRU	
Local Cache size (number of documents)	35	5-10-35-50
Redirection Cache size (number of registers)	35	
Simulation time (s)	20000	
Warm-up period (s)	4000	
Coverage radio (meters)	250	
MAC Protocol	802.11b	
Radio propagation model	Two Ray Ground	
Routing protocol	AODV	

The document request pattern follows a Zipf-like distribution, which has been demonstrated to properly characterize the popularity of the documents in the Internet [27]. The Zipf law asserts that the probability *P(i)* for the *i*-th most popular document to be requested is inversely proportional to its popularity ranking as shown in (1).

$$P(i) = \frac{\beta}{i^{\alpha}} \tag{1}$$

The parameter α is the slope of the log/log representation of the number of references to the documents as a function of its popularity rank (*i*). Values between 0.4 and 0.8 have been selected for the Zipf slope (with a default value of 0.8) in order to model low and high temporal locality.

Every node in the network implements a local cache that employs the LRU replacement policy. The default cache size allows storing 35 documents (35000 bytes). Cache sizes of 5, 10 and 50 documents have also been considered. In order to avoid cache misses due to the emptiness of the caches at the start of the simulation [28], a warm-up time has been considered using the first 20% of the simulation time. As the simulation time has been set to 20000 seconds, the warm-up time has a value of 4000 seconds. During the warm-up period, the performance metrics are not computed. Consequently, the analyzed statistics correspond to the time after the warm-up time. The redirection cache has the capacity of storing 35 registers.

As the coverage radio of the nodes is 250 meters and the simulation area is 1000x1000 m^2, the connectivity among neighbour nodes is different for each evaluated grid configuration. Figure 2 shows the connectivity for the 5x5, 7x7 and 9x9 grid configurations. As it can be observed, as the density of nodes increases, the number of neighbour nodes grows.

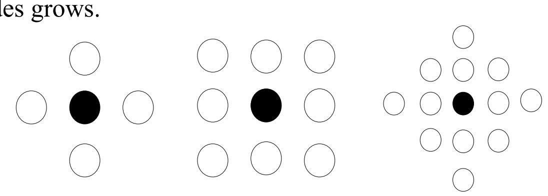

Figure 2. One hop connectivity of a node for 5x5, 7x7 and 9x9 grids.

The parameters employed in the rest of the evaluated caching schemed (HybridCache, DPIP, SimpleSearch and MobEye) are those proposed by their authors.

As performance metrics we consider:

- Traffic load: It measures the mean amount of traffic generated or forwarded by each node during the simulation. As the wireless medium is limited, the greater the generated traffic the greater the probability of interferences and collisions.
- Delay: It is defined as the mean time that a request requires to be served, that is to say, the mean time that a user will have to wait to receive the requested document.
- Timeouts: This metric defines the percentage of requests that have failed and have been requested again because the document has not been received before the timeout.
- Number of served documents: It is defined as the mean number of documents that have been retrieved to the nodes of the network during the simulation time. This metric shows the caching schemes that can serve more documents in a predefined time.
- Cache hits: They can be divided into local cache hits and remote cache hits. The local cache hit measures the percentage of requests served by the local cache. A local cache hit implies a zero delay, avoids the traffic generation in the network and reduces the servers load. The remote cache hit measures the percentage of requests served by network nodes different from the servers. A remote cache hit implies the reduction of the servers load. Depending on the cache scheme, the remote cache hit can be categorized into other classes. Taking into consideration the CLIR caching scheme, interception, cross-layer interception and redirection cache hits are considered as remote hits.
- Redirection hits: It is defined as the percentage of redirections that have been correctly resolved, i.e.,

the document has been served to the requester after a redirection. This metric measures the goodness of the redirection cache technique for those caching schemes that implement it.

The figures presented in this section correspond to the evaluation of a 7x7 grid network as the results obtained with the 5x5 and 9x9 networks are very similar. The performance evaluation will be studied as a function of the time between requests, the TTL of the documents, the Zipf slope and the local cache size. The redirection hit metric will be applied only to those caching schemes that implement the redirection technique, that is, CLIR, COOP and HybridCache.

B. *Time between requests*

Figure 3a represents the mean processed traffic by each node as a function of the time between requests. CLIR, DPIP and HybridCache are the caching schemes that generate the lowest traffic, followed by No Cache and SimpleSearch. MobEye generates more traffic because of the use of broadcast messages.

Figure 3b compares the mean delay of the requests and replies. CLIR is the caching scheme with the lowest delay. In fact, it is the only scheme that obtains a lower delay than the option of not using caches. SimpleSearch and MobEye employ a four request-reply messages method, and hence, they experience a greater delay and a greater traffic generation as previously observed. COOP has not been shown in this figure due to the high delay obtained. This behaviour is caused by the timeout needed to perform the direct request to the *DS* after the broadcast request has failed. DPIP also achieves a high delay due to the *DPIP_Timer* parameter that fixes a lower bound to the messages delay. Finally, HybridCache achieves a low performance for high loaded networks although this performance is improved as the traffic load is decreased. This fact is due to redirection loops caused by a wrong redirection management (Figure 3f). When time between requests increases, the information stored in the redirection table is obsolete related to the documents stored in the local caches as they are evicted from the local caches before the information can be considered obsolete. As the number of evictions in the local caches decreases the redirection cache is able to obtain more redirection hits because it only takes into account the TTL of the documents to delete the information of the redirection cache.

Figure 3c shows the mean percentage of timeouts per node. HybridCache obtains a high percentage of timeouts due to the bad redirection management as previously explained. Similarly, COOP presents the same behaviour as HybridCache because of the same reasons. Finally, the rest of the caching schemes obtain a percentage of timeouts close to zero. In fact, this should be the normal behaviour of the caching schemes as the servers are always available and it is always possible to create a route to them.

Figure 3d illustrates the mean number of documents received per node. All the compared caching schemes obtain a similar performance except COOP and, especially, HybridCache for times between requests lower than 25 seconds. For high loaded networks (5 or 10 seconds between requests), HybridCache achieves a very low performance, obtaining only a half of the documents compared to the rest of the caching schemes due to the timeouts (Figure 3c) and redirection errors (Figure 3f).

Figure 3e depicts the mean percentage of local and remote cache hits as a function of the mean time between requests. MobEye is the caching scheme that obtains more hits, although its performance decreases as the time between requests increases. However, the rest of the caching schemes show the opposite behaviour. In SimpleSearch, the percentage of remote hits is practically zero for high loaded networks. As the requests are performed very close in time, the broadcast method is not employed because the route to the servers is already created, and hence, the requests are sent directly to them.

Figure 3f represents the mean percentage of redirection hits as a function of the mean time between requests. The performance of COOP is reduced as the traffic load is decreased because of the reduction of the redirection cache hits due to incorrect redirections of requests to nodes that have removed the documents from their local caches. On the other hand, CLIR obtains a performance close to 100% for all the traffic loads. Finally, HybridCache shows a redirection cache hit close to zero except for very low loaded networks where the performance reaches 5%.

C. *TTL of the documents*

Figure 4a represents the mean traffic processed by each node as a function of the mean TTL of the documents. CLIR, DPIP and COOP generate less traffic than no Caching for all the studied TTLs. HybridCache is very sensitive to the TTL of the documents and, as the TTL is increased, the generated traffic also soars. This behaviour is due to the redirection cache, which only takes into account this parameter to delete the information in the redirection cache. Consequently, if a node evicts a document from its local cache, the nodes with information about the location of this document in their redirection caches will maintain incorrect data.

Figure 4b compares the mean delay as a function of the mean TTL of the documents. CLIR is the caching scheme that obtains the lowest delay. HybridCache, as shown in the previous study, is very sensitive to the TTL and the delay is highly increased as the TTL is incremented. The rest of the caching schemes obtain delays greater than the case of no Caching due to the four messages needed to obtain the document.

Figure 4c shows the evolution of the percentage of timeouts as a function of the TTL of the documents. COOP and HybridCache are the caching schemes with a percentage of timeouts greater than zero due to the previously commented reason. In fact, the percentage of timeouts is highly increased in HybridCache for TTLs greater than 2000 seconds.

Figure 4d illustrates the mean number of documents received per node as a function of the mean TTL of the documents. As it can be observed, a similar number of documents is obtained by all the caching schemes except COOP, with a slightly lower performance, and HybridCache with a drastically reduction of the retrieved documents with a mean TTL greater than 1000 seconds. This behaviour shows that HybridCache is very sensitive to the TTL of the documents because of a bad management of the redirection cache (as shown in Figure 4f). This figure reports that the redirection hits decrease until practically zero for a mean TTL of the documents greater than 1000 seconds. The incorrect management of the redirection cache also causes an increment of the traffic (Figure 4a), a greater number of timeouts (Figure 4b) and a higher delay (Figure 4c)

Figure 4e depicts the percentage of local and remote cache hits as a function of the mean TTL of the documents. As the mean TTL of the documents increases, they are stored for more time in the local caches; hence, the probability of a cache hit is also incremented. HybridCache obtains a good performance for low TTLs although this performance is drastically reduced when the documents do not expire. CLIR achieves a performance similar to the rest of the caching schemes and it is only overcome by MobEye and DPIP for high TTLs because of the use of broadcast message. On average, SimpleSearch is the caching scheme with the lowest cache hit ratio.

Finally, Figure 4f compares the percentage of redirection cache hits as a function of the TTL of the documents. As it was previously commented, HybridCache achieves a very high redirection hit rate for small TTLs. However, as the TTL increases, this rate drastically falls and reaches a percentage of zero for TTLs greater than 2000 seconds. On the other hand, it can be observed that COOP is also sensible with the TTL of the documents, as it obtains a lower percentage of redirection hits when they become obsolete more frequently. As the mean TTL increases, the redirection cache hit of COOP is also increased. Finally, CLIR achieves a percentage of redirection hits close to 100 % for all the studied TTLs.

D. Zipf slope

Figure 5a depicts the mean traffic processed by node as a function of the Zipf slope. CLIR is the caching scheme that obtains the lowest delay for all the slopes while MobEye and SimpleSearch generate more traffic than the No Caching option due to the broadcast requests. On the other hand, HybridCache also generates more traffic than the No Caching scheme for low slopes. This behavior is due to the replacement policy implemented by HybridCache, called SxO (Size x Order) [21]. This replacement policy is very sensitive to the popularity of the documents. Consequently, a low Zipf slope causes the reduction of the local cache hits, increasing the traffic generated in the network.

Figure 5b compares the mean delay as a function of the Zipf slope. The delay obtained by COOP is not shown because it is much greater than the rest of the caching schemes. Only CLIR and HybridCache (for a slope of 1.0) obtain a lower delay than the No Caching scheme. DPIP has a delay of even three times greater than CLIR although this difference is reduced as the Zipf slope increases. CLIR is the caching scheme with the lowest delay for all the considered Zipf slopes.

Figure 5c shows the mean percentage of timeouts per node as a function of the Zipf slope. As observed in previous studies, only HybridCache, COOP and MobEye present a percentage of timeouts different to zero. The behaviour of HybridCache and COOP is due to the incorrect implemented redirection technique. Nevertheless, the percentage of timeouts of these caching schemes is decremented as the Zipf slope increases because, as the Zipf slope increases, the percentage of local and remote cache hits increases and the documents can be served before the timeout. The rest of caching schemes obtain a percentage of timeouts close to zero.

Figure 5d illustrates the mean number of documents retrieved per nodes as a function of the Zipf slope. As observed in previous the studies, HybridCache achieves the lowest results, although this value is incremented as the Zipf slope is incremented because of the rise of the cache hits (Figure 5e). The errors produced by the redirection cache mechanism (Figure 5f) causes high occurrences of timeouts (Figure 5c), and hence, the reduction of the number of obtained documents. Similarly, COOP also achieves a lower performance than the rest of caching schemes. CLIR, DPIP, SimpleSearch and MobEye obtain a similar number of documents for all the studied Zipf slopes.

Figure 5e compares the percentage of local and remote cache hits as a function of the Zipf slope. The Zipf slope indicates the probability that the more popular documents will be requested more times. Hence, if the Zipf slope increases, the percentage of local cache hits is expected to increase too. This behaviour is confirmed in this figure. MobEye, as occurred in the previous studies, is the caching scheme that achieves a higher cache hit percentage because of the use of broadcast requests. CLIR presents a performance similar to the rest of the caching schemes.

Figure 5f depicts the percentage of redirection cache hits as a function of the Zipf slope. As observed in the previous studies, HybridCache obtains a redirection cache rate close to 0 %. COOP reduces the percentage of cache hits as the Zipf slope increases. On the contrary, CLIR achieves the same performance (close to 100 %) for all the studied values of the Zipf slope.

E. Cache size

Figure 6a depicts the mean processed traffic by the nodes as a function of the local cache size. As the cache size rises the generated traffic is decreased because the probability of a local cache hit is increased. CLIR, DPIP and COOP are the caching schemes that generate a lower traffic than the No Caching scheme for all the studied cache sizes. MobEye is the caching scheme that generates more traffic due to the use of broadcast requests. On the other hand, HybridCache only performs better than No Caching when the cache size is greater than 20 documents. Hence, HybridCache does not work correctly when using small caches due to the implemented SxO replacement policy.

Figure 6b compares the mean delay as a function of the local cache size. CLIR is the caching scheme with the lowest delay and, in this case, is the one that performs better than the No Caching scheme for all the studied cache sizes. HybridCache presents a big delay for small caches, although it is drastically reduced as the cache size increases. In addition, SimpleSearch and MobEye always obtain a bigger delay than the No Caching scheme for all the studied cache sizes due to the four messages needed to obtain a document. Finally, DPIP shows a delay close to 150 milliseconds due to the limit imposed by the *DPIP_Timer*.

Figure 6c presents the mean percentage of timeouts as a function of the local cache size. As observed in previous studies, only HybridCache, MobEye and COOP show a percentage of timeouts different to zero. This percentage is reduced, especially in HybridCache, as the cache size increases because the probability of local and remote cache also augments.

Figure 6d shows the mean number of documents retrieved per node as a function of the cache size. HybridCache is very sensitive to the local cache size as it obtains the fewest documents with small caches. However, this number of retrieved documents is incremented as the local cache size is augmented. This behaviour is caused by the SxO replacement policy as it does not work appropriately for small caches as it obtains poor cache hits (Figure 6e). On the other hand, COOP also obtains fewer documents than the rest of caching schemes. Finally, we can remark that the number of documents retrieved in the simulation time by CLIR, DPIP and SimpleSearch is not dependent on the cache size.

Figure 6e illustrates the mean percentage of cache hits as a function of the cache size. The greater the cache size is, the higher the percentage of cache hits is expected to obtain as can be observed in the figure. MobEye is the caching scheme with the higher cache hit percentage, especially the remote cache hit because of the broadcast requests. CLIR achieves, as well as MobEye and SimpleSearch, the best local cache hits.

Figure 6f compares the percentage of redirection hits as a function of the cache size. The shown behaviour is similar to that of previous studies as HybridCache obtain a performance close to 0%. COOP achieves about 85% while CLIR obtains a performance between 95% and 100%. However, there is a great variability with small cache sizes due to the great probability of replacements.

V. CONCLUSIONS

In this paper, we have evaluated the performance of the CLIR caching scheme applied to static grid ad hoc networks. This evaluation has been performed using the following metrics: mean traffic processed by the node, the delay perceived to obtain the requested documents, the mean percentage of timeouts, the mean number of retrieved documents per node, the local and remote cache hits and the percentage of redirection cache hits. We have evaluated the influence of the traffic load in the network, the TTL of the documents, the traffic pattern (Zipf slope) and the local cache size. In addition, we have compared the performance of CLIR to the caching schemes HybridCache, COOP, DPIP, SimpleSearch and MobEye. Finally, the performance of CLIR has also been compared to the case of an ad hoc network that does not implement any caching scheme.

From the set of developed simulations we can conclude that MobEye, COOP and HybridCache are not suitable for static ad hoc networks. We base this conclusion on the fact that they obtain a mean percentage of timeouts different to zero. This behaviour is not acceptable in this kind of networks where the servers are always available because the wireless nodes do not move and hence, they are always available. On the other hand, they also retrieve fewer documents than the rest of the caching schemes. Taking into account the rest of caching schemes (DPIP, SimpleSearch and CLIR), CLIR always obtains the lowest traffic generation as well as the lowest delay for all the studied situations. In addition, CLIR always presents a better performance than the No Caching Scheme for all the studied parameters and, hence, we can assert that it is suitable for this kind of networks.

Finally, CLIR presents a percentage of local cache hits similar to the rest of caching schemes and it obtains the best percentage of redirection cache hits, with a performance close to 100%. This result demonstrates that the management of the redirection cache is efficient.

As a future work we propose to evaluate the influence of the employed routing protocol on the caching scheme performance. As we have developed the piggy-backing method of CLIR using OADV, different behaviours are expected using other routing protocols.

ACKNOWLEDGEMENTS

We would like to thank Adela Isabel Fernandez-Anta for revising the syntax and grammar of this article. This study was partially supported by the National Project No. TEC2009-13763-C02-01.

REFERENCES

[1] F. J. González-Cañete and E. Casilari, "Evaluation of a Cooperative Caching Scheme for Grid Ad Hoc Networks", Proc. 6th International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies (UBICOMM 2012), 2012, pp. 97-103.

[2] D. Wessels, Web Caching: Reducing Network Traffic. O'Reilly, 2001.

[3] P. Kuppusamy and K. Thirunavukkarasu, B.Kalaavathi, "A Review of Cooperative Caching Strategies in Mobile Ad Hoc

Networks", International Journal of Computer Applications, vol. 29, no. 11, 2011, pp. 22-26.

[4] G. Dodero and V. Gianuzzi, "Saving Energy and Reducing Latency in MANET File Access", Proc. 26th International Conference on Distributed Computing Systems Workshops (ICDCSW'06), 2006, pp. 16-20.

[5] S. Lim, W.C. Lee, G. Cao and C.R. Das, "A novel caching scheme for improving Internet-based mobile ad hoc networks performance", Ad Hoc Networks, vol. 4, no. 2, 2006, pp. 225-239.

[6] S. Lim, W.C. Lee, G. Cao and C.R. Das, "Cache invalidation strategies for Internet-based mobile ad hoc networks", Computer Communications, vol. 30, no. 8, 2007, pp. 1854-1869.

[7] T. Moriya and H. Aida, "Cache Data Access System in Ad Hoc Networks", Proc. 57th IEEE Semiannual Vehicular Technology Conference (VTC 2003), April 2006, vol. 2, pp. 1228-1232.

[8] B. Tang, H. Gupta and S.R. Das, "Benefit-Based Data Caching in Ad Hoc Networks", IEEE Transactions on Mobile Computing, vol. 7, no. 3, 2008, pp. 289-304.

[9] Y.H. Wang, J. Chen, C.F. Chao and C.C. Chuang, "A Distributed Data Caching Framework for Mobile Ad Hoc Networks", Proc. 2006 International conference on Wireless communications and mobile computing, 2006, pp. 1357-1362.

[10] J. Cho, S. Oh, J. Kim, K.H. Lee and J. Lee, "Neighbor Caching in Multi-Hop Wireless Ad Hoc Networks", IEEE Communications Letters, vol. 7, no. 11, 2003, pp. 525-527.

[11] P. Nuggehalli, V. Srinivasan and C.F. Chiasserini, "Energy-Efficient Caching Strategies in Ad Hoc Wireless Networks", Proc. 4th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2003), June 2003, pp. 25-34.

[12] N. Chand, R.C. Joshi and M. Misra, "Cooperative Caching in Mobile Ad Hoc Networks Based on Clusters", International Journal on Wireless Personal Communications, no. 43, 2007, pp. 41-63.

[13] M.K. Denko, "Cooperative Data Caching and Prefetching in Wireless Ad Hoc Networks", International Journal of Business Data Communications and Networking, vol. 3, no. 1, 2007, pp. 1-15.

[14] V. Gianuzzi, "File Distribution and Caching in MANET", Technical Report DISI-TR-03-03, DISI Tech University of Genova (Italy), 2003.

[15] Y. Du and S. Gupta, "COOP – A cooperative caching service in MANETs", Proc. Joint International Conference on Autonomic and Autonomous Systems and International Conference on Networking and Services (ICAS-ICNS 2005), October 2005, pp. 58-63.

[16] A. Klemm, C. Lindemann and P.D. Waldhorst, "A Special-Purpose Peer-to-Peer Sharing System for Mobile Ad Hoc Networks", Proc. IEEE Semiannual Vehicular Technology Conference (VTC 2003), October 2003, pp. 2758-2763.

[17] G. Chiu and C. Young, "Exploiting In-Zone Broadcast for Cache Sharing in Mobile Ad Hoc Networks", IEEE Transactions on Mobile Computing, vol. 8, no. 3, 2009, pp 384-397.

[18] C.Y. Chow, H.V. Leong and A. Chan, "Peer-to-Peer Cooperative Caching in Mobile Environments", Proc. 24th International Conference on Distributed Computing Systems Workshops (ICDCSW'04), March 2004, pp. 528-533.

[19] H. Artail, H. Safa, K. Mershad, Z. Abou-Atme and N. Sulieman, "COACS: A Cooperative and Adaptive Caching Systems for MANETs", IEEE Transactions on Mobile Computing, vol. 7, no. 8, 2008, pp. 961-977.

[20] C.Y. Chow and H.V. Leong, A. Chan, Group-based Cooperative Cache Management for Mobile Clients in a Mobile Environment, Proceedings of the 33rd International Conference on Parallel Processing (ICPP'04), 2004, pp. 83-90.

[21] L. Yin and G. Cao, "Supporting Cooperative Caching in Ad Hoc Networks", IEEE Transaction on Mobile Computing, vol. 5, no. 1, 2006, pp. 77- 89.

[22] Y. Ting and Y. Chang, "A Novel Cooperative Caching Scheme for Wireless Ad Hoc Networks: GroupCaching", Proc. International Conference on Networking, Architecture and Storage (NAS 2007), 2007, pp. 62-68.

[23] N. Chand, R.C. Joshi and M. Misra, "Efficient Cooperative Caching in Ad Hoc Networks", Proc. 1st International Conference on Communication System Software and Middleware (Comsware'06), January 2006.

[24] F. Sailhan and V. Issarny, "Cooperative Caching in ad hoc Networks", Proc. 4th ACM International Conference on Mobile Data Management (MDM'2003), January 2003, pp. 13-28.

[25] F.J. González-Cañete, E. Casilari and A. Triviño-Cabrera, "A cross layer interception and redirection cooperative caching scheme for MANETs", EURASIP Journal on Wireless Communications and Networking 2012, 2012:63, doi:10.1186/1687-1499-2012-63.

[26] http://www.isi.edu/nsnam/ns/ [retrieved: February, 2013]

[27] L.A. Adamic and B.A. Huberman, "Zipf's law and the Internet", Glottometrics, vol. 3, 2002, pp. 143-150.

[28] S.G. Dykes, K.A. Robbins and C.L. Jeffery, "Uncacheable Documents and Cold Starts in Web Proxy Cache Simulations", Technical Report CS-2001-01, Texas University (EEUU), 2000.

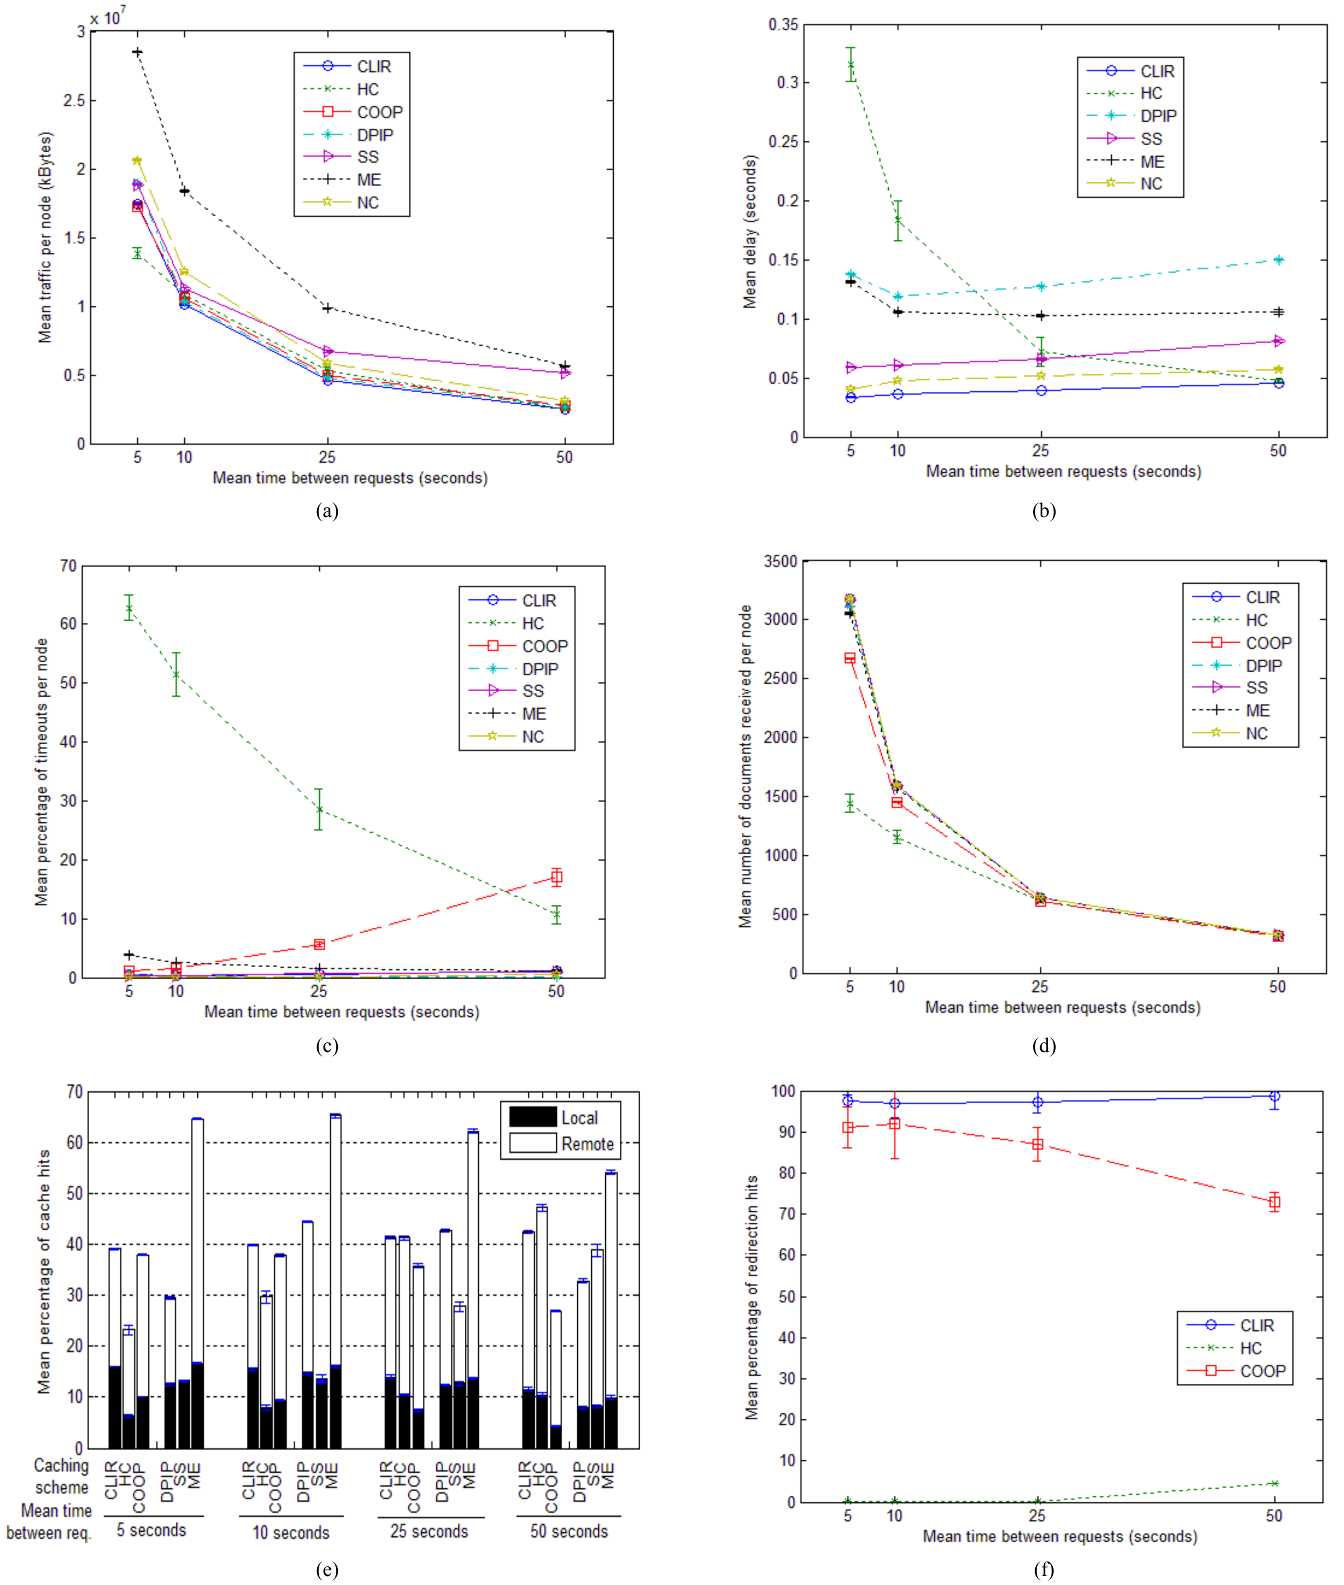

Figure 3. Mean traffic processed by node (a), delay (b), percentage of timeouts (c), documents received (d), cache hits (e) and redirection cache hits (f) as a function of the mean time between requests.

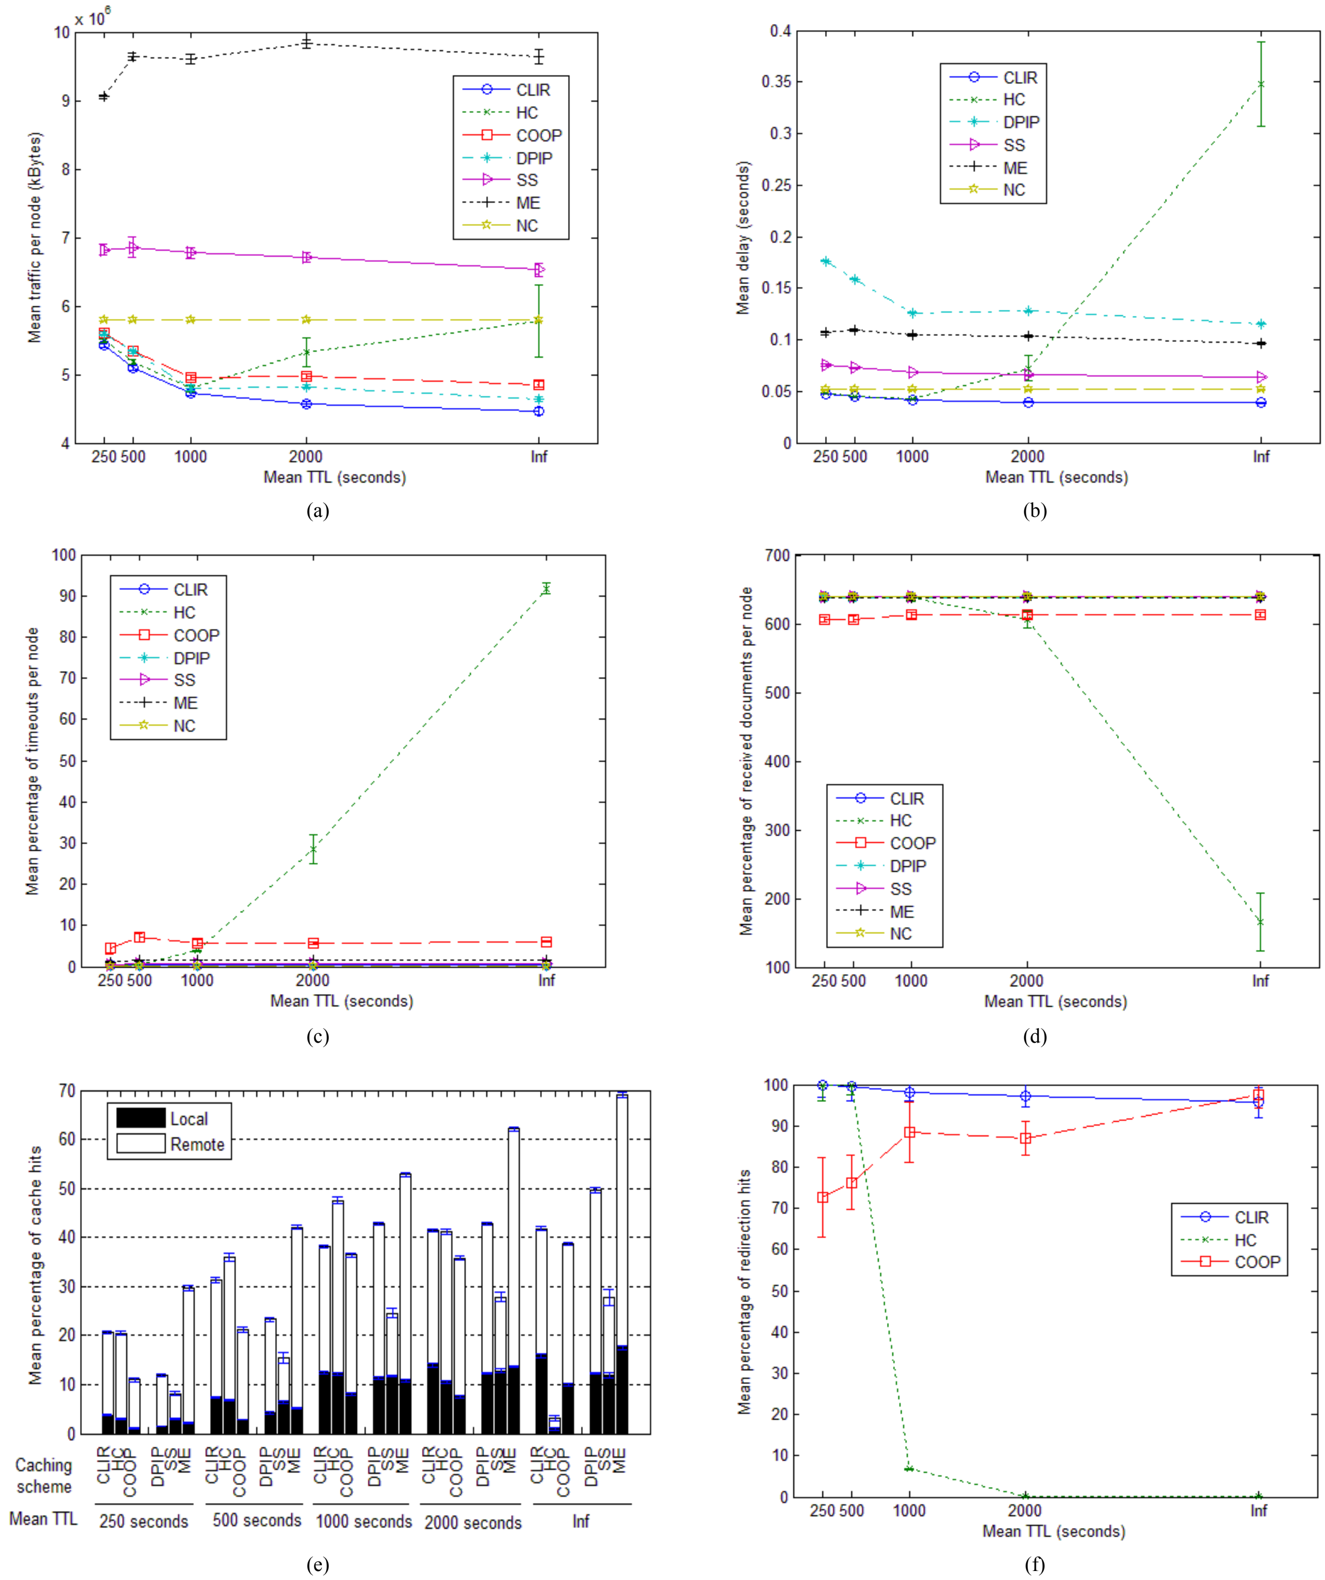

Figure 4. Mean traffic processed by node (a), delay (b), percentage of timeouts (c), documents received (d), cache hits (e) and redirection cache hits (f) as a function of the mean TTL of the documents.

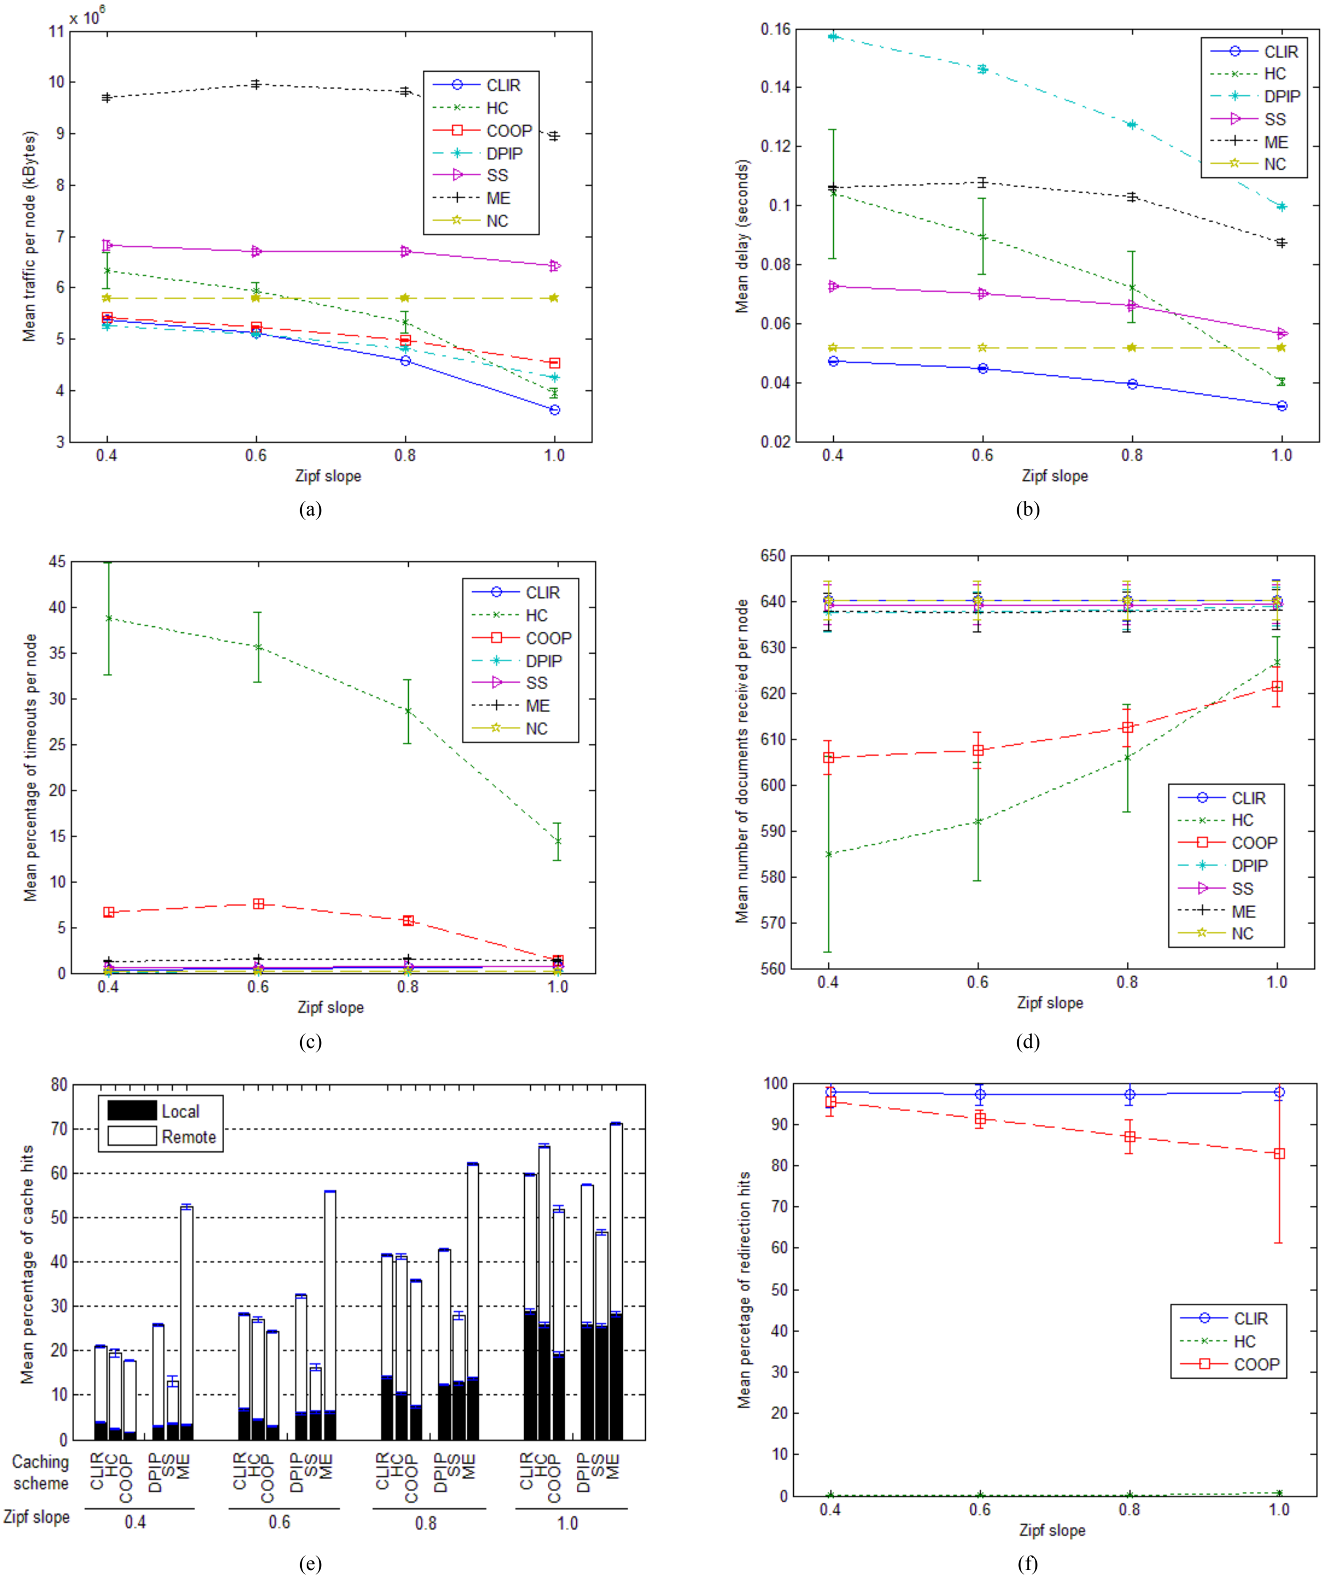

Figure 5. Mean traffic processed by node (a), delay (b), percentage of timeouts (c), documents received (d), cache hits (e) and redirection cache hits (f) as a function of the Zipf slope.

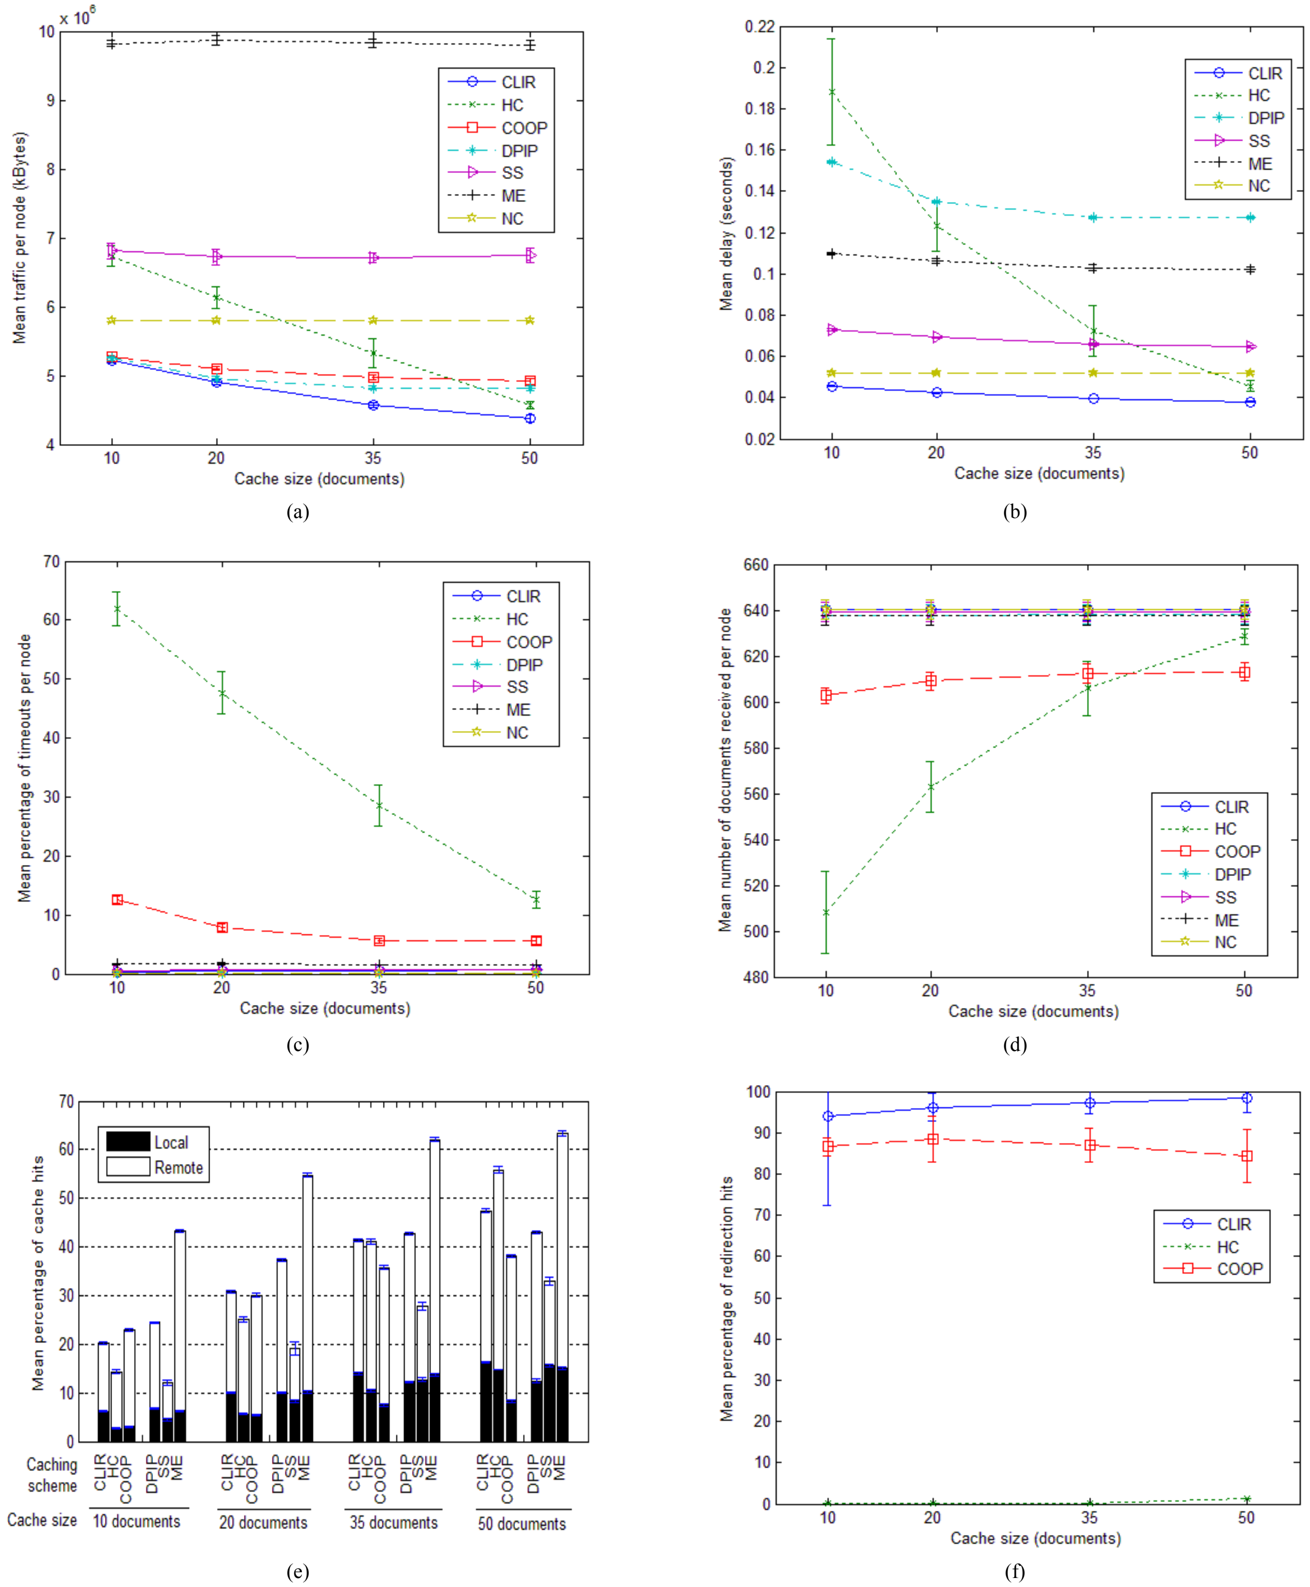

Figure 6. Mean traffic processed by node (a), delay (b), percentage of timeouts (c), documents received (d), cache hits (e) and redirection cache hits (f) as a function of the cache size.