Maximising Access to a Spectrum Commons using Interference Temperature Constraints

Joe Bater*, Hwee-Pink Tan†, Kenneth N Brown*, Linda Doyle†
Centre for Telecommunications Value-chain Research,
*Cork Constraint Computation Centre, Dept. of Computer Science, University College Cork, Ireland.
Email: j.bater@4c.ucc.ie, k.brown@4c.ucc.ie
†Dept. Electronic Engineering, Trinity College Dublin, Ireland,
Email: tanhp@tcd.ie, ledoyle@tcd.ie

Abstract— **We propose a new spectrum-access etiquette for cognitive radios in a spectrum commons. When congestion might block a new device from initiating a call, the surrounding devices coordinate their actions and locally reassign their spectrum to create a gap for the new entrant. The etiquette is designed for devices operating dissimilar services with different bandwidth and quality requirements. It generates link-level interference temperature constraints and finds a satisfying assignment using local search. In experimental simulation, we demonstrate that the etiquette provides significantly higher completion rates while improving the quality of the completed calls.**

I. INTRODUCTION

The driving force for new radio spectrum regimes has been the inefficiency of existing mechanisms. The most basic of these range from "spectrum licensing" providing exclusive use for a given contractual period, to having a "spectrum commons" [1], [2]. A spectrum commons (e.g. [3]) is basically an allotment of unlicensed spectrum which any user may access opportunistically. However, a completely unregulated commons is unlikely to be successful, and so some rules and etiquettes for access and behaviour are likely to be required. The key question is then: can we devise an etiquette which is both desirable and technically feasible?

To get maximum spectrum usage, it is likely that many different types of traffic, with widely different bandwidth, power and range requirements, will have to co-exist in the commons [3], [4]. While there is increasing interest in etiquettes which can work with dissimilar or heterogeneous services, e.g. 802.11a/b/g, 802.16, Bluetooth and UWB sharing the spectrum "horizontally" (having equal priority and status), most proposed (and implemented) etiquettes only consider homogeneous sets of devices/services. These etiquettes are usually described by whether they are cooperative (or not), or pro-active (rather than reactive). Cooperative etiquettes may be centralised - depending on a server, e.g. [5]; or distributed (often involving "cognitive radios" to sense the environment, make decisions and shape their transmitted signal - in terms of output power, frequency mask, codebook etc.) e.g. [6].

Some etiquettes involve "simple" operations conducted by an individual radio independently of other radios to improve global utilisation of radio resources. More complex etiquettes require interaction amongst devices to locally optimise their use of the spectrum. These might rely on device priority or micro-auctions to determine which device might pay the most for the spectrum and assign that spectrum accordingly e.g. [7], or redress the assignments locally e.g. [8]. Regardless of the mechanism, these etiquettes usually require a means of communication and a "common signalling control channel" (CSCC) is widely proposed [3], [7].

We propose and implement a local reassignment etiquette for dissimilar wireless systems, using link-level interference temperature constraints which faithfully capture the service requirements of the links involved. The aim of this etiquette is to increase access by reassigning spectrum to produce a "spectrum hole" which allows additional call sessions. We demonstrate in an experimental simulation that this etiquette offers a significant increase in the number of completed calls, and simultaneously improves the quality of the completed calls.

The remainder of this paper is structured as follows: In Section II, we describe our assumptions regarding the spectrum commons. In Section III we present the interference temperature (IT) constraints which capture the requirements of the services used. In Section IV, we show how to use these constraints for local spectrum access in a heterogeneous traffic model. Finally, in Section V, we describe an experimental simulation of this new etiquette in comparison to "standard" etiquettes.

II. A SPECTRUM COMMONS FOR HETEREGENEOUS SERVICES

We propose an environment, similar to [3], in which a number of reconfigurable devices capable of delivering dissimilar services access a spectrum commons in a co-operative manner. We assume that the spectrum is divided into discrete non-overlapping and non-interfering channels in the frequency domain, where every channel is available to all users.

We consider a distributed but static deployment of reconfigurable transmitting/receiving devices across a universal geographic area (for example, a large office environment). Each device is capable of reconfiguring itself to provide basic services that differ in required service level (minimum receive power), bandwidth (the number of channels that must be contiguously assigned) and threshold carrier-to-interference ratio (C-I) for error tolerant operation. The C-I determining the call quality is given by the ratio of the desired signal power to the sum of received unwanted co-channel signal powers (plus the noise floor).

At any time, an idle device may attempt to initiate a (unicast) call to another device on a particular service. If the other device is also idle, and the call is allowed by the etiquette, the transmitting device sets its transmit power to the minimum power needed to achieve the required service level. After the call is completed, the spectrum is released.

We assume that there is a common signalling control channel (CSCC) on which devices report their state (IDLE, BUSY-TRANSMITTING, BUSY-RECEIVING), the service identity, transmit power, and their channel(s). The CSCC can be used as a beacon indicating the signal loss between devices. It will also be used to pass messages relating to the reassignment protocol described here and to obviate the hidden-transmitter problem.

III. IT SPECIFIC CONSTRAINTS

Let us assume that device t wishes to initiate a service s call to a nearby device r. Interference *constraints* are used to determine the channel assignments for (t, r) as well as other devices in their locality so that interference remains within levels acceptable to service s. A constraint consists of a *scope* (a subset of variables in the problem), and a *relation* (a set of simultaneously permitted/disallowed value assignments for the variables in scope). If the scope constrains two variables *only* then the constraint is termed a *binary* constraint; if several (more than two) variables are in scope then the constraint is termed *non-binary*.

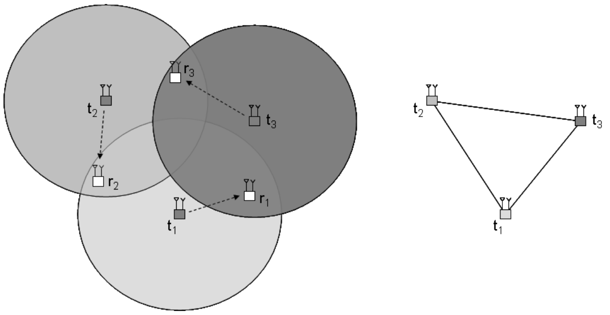

Fig. 1. Translation of the radio model into binary Graph Colouring constraints.

The most common constraints in spectrum utilisation literature are *binary* since they are derived from Graph Colouring models, e.g. [5], [8]. These are specified in terms of a *re-use distance* within which no pair of devices should be permitted to simultaneously transmit on the same channel - see Fig. 1. The detectable signal power of the transmitting devices (shown as dark grey devices) radiates omni-directionally and falls off quickly with distance. The output power of transmitting devices determines the radius of these "signal discs", and is dictated by the particular service level required of the receiving device (or devices). However, to another set of communicating devices that signal acts as interference, transmitter pairs are typically constrained so that the "discs" may not overlap, within this distance devices should not operate on the same spectrum.

Rather than simply forbidding all interference, the FCC's Spectrum Policy Task Force propose an Interference Temperature (IT) metric [9], [10], as a measure of interference power that could be received without causing excess interference. The IT model moves away from the most common transmitter-centric approach towards one which takes into account all unwanted signals in the RF environment, and from these calculates an upper bound on the signals that could be introduced into the system to enable additional communications. The ideal IT model would occur on a link-by-link basis, though a generalised model of IT providing a single limit on permissible interference has also been suggested [11] (see Fig. 2).

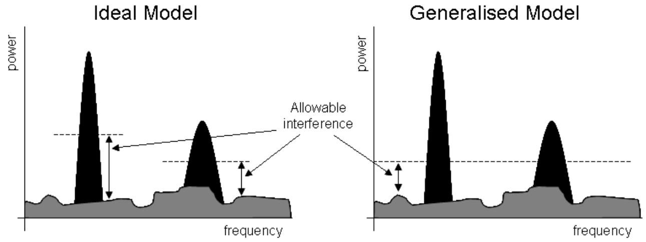

Fig. 2. Ideal and Generalised Interference Temperature (from [11])

In Fig. 2 we see two signals (the dark areas depict the power spectra in tuned frequency channels) being operated in a noisy frequency band (the light grey area represents the noise power spectra across the entire channel allotment). The noise spectra is irregular across the band largely due to environmental effects. Both signals are currently operating in the absence of unwanted signal interferers, consequently they may yet receive further interference before their required C-I levels reduce to the minimum required for error tolerant service. These levels are shown in the left sub-figure for a link-by-link basis, and a generalised basis on the right.

In [12], we show that the binary graph-colouring constraints are inefficient in terms of spectrum usage - that is, they deny assignments that would give acceptable call quality. Instead, we propose constraint models that are in-line with the IT metric by considering the maximum allowed interference within the bandwidth of the call (rather than as a general measure of "interference" across a wide spectrum band). Consider a receiver R receiving a signal from a transmitter T with power P. There is a group G of potential interferers for R (i.e. transmitters broadcasting such that their signal is received by R at a known power, and which could cause some interference were they are co-channel with T). The task is then to assign channels to T and G such that R's C-I is greater than the defined service level. That is, $P/(N + \sum P_i) > \theta$, where the sum is over all co-channel transmitters T_i in G, P_i is the power received by R from T_i, N is the background noise, and θ is the required C-I threshold. This gives a limit as to the maximum amount of interference tolerable at the receiver in the bandwidth of the specific communication, which we set as our *interference temperature limit*.

Rather than implement this constraint as written, we compile it down to a list of maximal sets of co-channel transmitters: any channel assignment such that the set of co-channel transmitters is a subset of one of those in the list is then a valid assignment. Our list defines a relation over the set of potential interferers. Each element of the relation is a tuple of 0s and 1s, where 0 denotes that the transmitter can be co-channel.

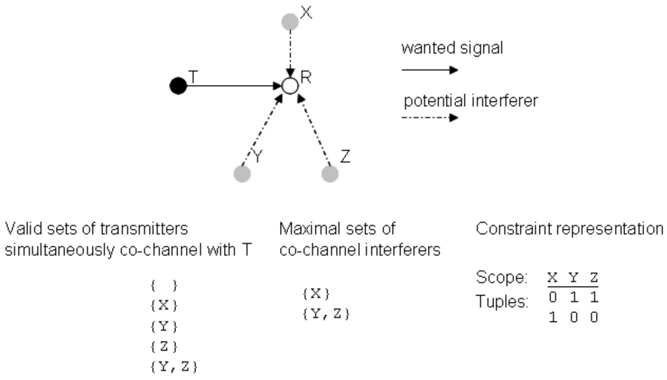

Fig. 3. Tuples in a Non Binary Constraint

Consider the situation in Fig. 3, where a transmitter T is sending a signal to receiver R. There are three other active transmitters (X, Y and Z) whose signals could interfere with R. The task is to specify what subsets of {X, Y, Z} can be simultaneously co-channel with T, such that the interference received by R is sufficiently low to allow its call to proceed with acceptable quality. Suppose now that X, Y and Z are transmitting at power levels such that any one of them on their own can be co-channel with T, or that Y and Z together can be co-channel with T, but that X and Y together cannot be simultaneously co-channel, and neither can X and Z be simultaneously cochannel with T. Clearly, if none of X, Y or Z are co-channel with T, then the call is acceptable. The list of acceptable co-channel sets is shown in column (i) of Fig. 3 From this, we obtain a smaller list by representing only the maximal sets, such that any subset of a maximal set is acceptable (column (ii)). Finally, column (iii) shows the compiled constraint representation of these maximal sets. In practice, there can be arbitrarily many devices in the scope of the constraint, and a large number of table entries.

IV. SPECTRUM ETIQUETTE USING IT CONSTRAINTS

The key with any multi-agent reasoning system (including Cognitive Radio) is to determine its strategic reasoning capability, in this case that which allows devices to "negotiate" shifts in existing spectrum use. Local reassignment (bargaining) etiquettes have previously been proposed in [8] and [13]. In [8] devices in a mobile ad-hoc network (with continuously backlogged traffic) are reduced to nodes in a GRAPH COLOURING problem, and are to simultaneously transmit using as many colours as possible. Triggered by mobility events forcing changes in the underlying constraint graph, nodes bargain to exchange colours if and as necessary when computing reassignments. The complexity of this locally bargained spectrum reassignment is significantly less than recolouring the entire graph with little, if any, loss of optimality in the assignment. Chen et al. [13] apply the same approach in a multi-cellular wireless radio access network (WRAN) system operating in digital television (DTV) spectrum. Whereas [8] aims to maximise system fairness, [13] aims to reduce blocking probability in the face of inter- and intra-cellular binary interference constraints.

The aim of our etiquette is to improve overall spectrum utilisation (access to spectrum) whilst ensuring that the quality of existing calls is not adversely affected. That is, when a new call is requested, its impact is assessed; if no unacceptable interference would be created, the call proceeds; if it would cause unacceptable interference, then rather than block it we attempt to reassign the channels of neighbouring devices to create a spectrum "hole" in which we can place the call; if no such reassignment is found, then the call is refused.

Initially a intending-to-transmit device T broadcasts on the CSCC its intent to make a new communication. The intended receiver R (and any other currently receiving devices who can hear the new transmitter) generate constraints taking the potential new transmission as well as existing transmissions from potential interferers to R into account. We will call this the *core problem*. This problem constrains the spectrum assignments available to T so that R might achieve an adequate C-I call and that T will not inflict too much interference on existing communications. In the case that the core problem is solvable (i.e. assignments for existing devices do not need to be altered to fit a new assignment at T), then the new transmitter sets up a call on a suitably selected block of spectrum.

Otherwise, a *reassignment group* (RG) is formed. Here the transmitters in scope of the core problem constraints may be directed to change their assignments in order to accomodate the new call, provided that they can do so without harmfully affecting their own and other ongoing calls. *Peripheral* constraints now need to be generated on the receivers of the reassignment group devices. These restrict the reassignments permitted to transmitters in the RG due to the interference effects of and on calls beyond the range of the intended call. If an RG transmitter is to change assigned channel(s), then it must not contribute excessive interference to other nodes (nor degrade its own call quality further than that demanded by the service involved).

If the intended transmitter can find an assignment to the RG devices which simultaneously solves the core problem and peripheral problem constraints, then the new transmitter can require the transmitters in the RG to reassign themselves accordingly. Such an etiquette should have two advantages: (i) no call is placed if it causes unacceptable interference; and (ii) spectrum re-use is increased by reallocating channel assignments. We place the responsibility for following the etiquette on the initiating device, using information provided to it by the neighbouring devices, and the initiating device can use any method for computing a reassignment, as long as it guarantees to satisfy the constraints.

There are many methods in which the new transmitter can try to find an acceptable reassignment to the RG transmitters (solving all constraints). These methods can be categorised as either *complete* or *heuristic*, depending on the extent to which the search space (of possible spectrum reassignments) is explored. Any search algorithm which produces a reassignment satisfying the constraints would do. As long as the constraints are sent to the initiating node, then our proposed etiquette does not need to specify more - we leave it up to the initiating device to work out a satisfactory reassignment any way it chooses. In fact, each device could use its own algorithm. Low memory/low cpu devices could choose to do no reassignment, and simply abort if there is no original hole; other devices could try a cheap local search or randomised generate-and-test; or they could try more sophisticated local searches; or they could try a complete search.

V. EXPERIMENTAL EVALUATION

We compare our IT-based etiquette against three other etiquettes in a software simulation of a spectrum commons. Below we describe the model of the commons, the etiquettes and the experimental results. We hope to see more calls set up and completed (that is, not suffering unacceptable levels of interference at any point in the calls duration).

A. *"Commons Model"*

We simulate a spectrum commons in which 100 devices are randomly scattered (with a uniform distribution) across a 100m by 50m area, representing a large office-like environment. Each device is represented as an object, which maintains details of its own location (x,y coordinates) and communications. A CSCC object encapsulated in each device determines the predicted loss between devices, broadcasts its transmit power, busy state, operating channel(s), etc. This information is used by each device when called to apply the selected etiquette. These devices can opportunistically access 32 channels of spectrum modelled to be in the ISM (2.4GHz) band. Each device may select any of the various services outlined in Table I, all of which must operate in the same shared allotment. Services A and B might be voice and data services operating over very short range, e.g. cordless telephony etc., whereas Service C might be a data service operating over a wider bandwidth, and Service D a direct sequence spread spectrum service (though a relatively poor processing gain is assumed, hence a C-I target of only -3dB). These services might have different modulation schemes and/or codebooks are are completely undecodable to other devices.[1]

TABLE I

SPECIFICATION OF REQUIREMENTS FOR DISSIMILAR SERVICES

Service	No. of Channels	Service-Level (dBm)	Target C-I (dB)
A	1	-50.0	12.0
B	2	-55.0	9.0
C	4	-60.0	5.0
D	10	-70.0	-3.0

Rather than nodes continuously transmitting completely backlogged traffic, we consider a more dynamic scenario in which calls each having a random duration in the range [0, 10) determined at their outset are attempted. In this model simulations consisting of a specified number of iterations/time slots are run in which for every time slot the following process is repeated : (i) all calls whose duration has come to an end are terminated, the spectrum they were using released, and the devices involved revert to an idle-state; (ii) a randomised round-robin process selects each device which, if idle, can attempt to initiate a call to another idle device on a randomly selected service (provided the received signal achieves the minimum service level necessary to operate that service). Each device attempting to initiate a call does so probabilistically (if $x \geq y$ where x and y are random values in the range (0, 100)), otherwise remaining idle. If the outcome of the etiquette being applied is a success, then the device begins to transmit at the minimum power necessary to achieve their calls, subject to a maximum transmit power of 20dBm, otherwise the call is considered blocked. Any transmission is potentially seen as interference by other devices. We use a multi-wall propagation model based on COST-231 due to Tamminen [14], verified in the 2.4GHz band by [15], to determine the powers received by all other devices (though any other propagation model, e.g. ITU-R P.1238, could be applied). Finally, (iii) the C-I on each ongoing call is recalculated.

B. *Etiquettes*

Our control study, *Random*, is with devices randomly assigning their channels in a completely uncoordinated way. At call set-up each device simply selects a block of spectrum with no consideration to itself or other devices.

Listen Before Talk (LBT) schemes form the basis of most non-cooperative spectrum access schemes [16]. We consider two LBT variants, *LBT-simple* and *LBT-DFS*. LBT-simple

---

[1] In fact all parameters in our simulation are defined in a parameter file given to the simulation tool at runtime. Thus these can be changed easily, from the number of devices to the dimensions of the region considered. Our services are intended to be arbitrary and abstract, allowing the user of the tool to experiment with as wide a range of possible dissimilar service models. For example we could add a fifth service E, simulating DS-WCDMA operating over the full spectrum band of 32 channels (as also posited in the same input parameter file), with a service level only marginally above the noise floor of say -79dBm, with an effective processing gain the C-I requirement might be as low as -18dB (and a sixth F...); or alternatively reduce the number of anticipated services to just two which might be highly interference intolerant, the impact of which would be to reduce the interference temperature limit on the channels involved which in turn means that fewer possible interferers would be likely acceptable.

requires that a device wishing to transmit first scan the spectrum using a receiver of its own, it may then only select channels on which no interference can be detected. This aims to ensure an “interference free” (other than that from the noise floor) assignment at the intended receiver. The LBT-DFS variant is more like a dynamic frequency selection (DFS) scheme, in which we assume that there are insufficient “free” channels. The originating devices will scan their environment and proceed by selecting channels with only acceptable levels of interference [17].

Our proposed *Local Reassignment* (LR) etiquette using interference temperature constraints to accurately model the conditions at the receivers involved (as described in Section IV). We experimented with several reassignment algorithms including: *chronological backtracking*, a simple *generate-and-test* and a more systematic *iterated local search*, before settling on an greedy GRAPH COLOURING heuristic adapted from [5], which presented the best trade-off of speed and performance for the approaches investigated.

The *nmsb* algorithm presented by Peng *et al.* colours a graph by iteratively applying a labelling function and a colouring function. The *labelling* function calculates the “impact” of each uncoloured node on the remaining uncoloured nodes of a graph (e.g. an iterated max-degree heuristic) and returns that having the highest impact value, this node is then given to the *colouring* function which assigns it as best able. Once a node is assigned then it, and its incident edges are removed from the graph, until the graph is eliminated.

Since our constraints exist on *hyperedges* rather than simple edges, we adapt this algorithm by flagging a node as “assigned” or “unassigned”. As before, the labelling function determines the unassigned node with the highest adjacency count (of unassigned nodes) and labels it with a value which does not violate any constraints (constraints with partially unassigned scopes can still detect if the current partial assignment will lead to failure). Ties in the labelling function are broken arbitrarily, and as before the process repeats until all nodes to be reassigned have been reassigned.

Neither the original method, nor its adapted version, are guaranteed to provide an optimal result, but in practice very quickly results in relatively good colourings.

C. Results

Measurements taken for the period are shown in Fig. 4, in which we measure:

- the total number of call attempts
- the total number of call setups
- the total number of calls completed (the number of calls which had satisfactory C-I for the duration of the call)

Measurements are taken at each iteration for:

- the number of calls in session (as shown in Fig. 5),
- the average C-I of calls (by each service - as shown in Fig. 6),
- the relative number of calls (by each service) currently having an adequate C-I for the service (as shown in Fig. 7),

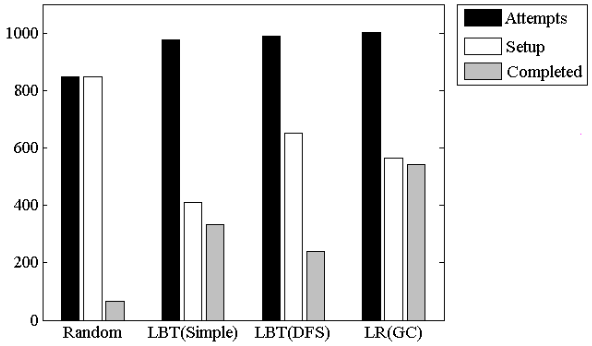

Fig. 4. Summary of calls attempted, setup and completed

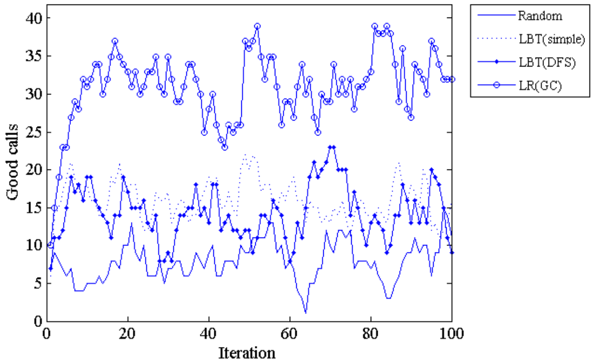

Fig. 5. Number of successful calls for different etiquettes

The results for Random show the disadvantages of uncoordinated access (if such a reminder was necessary): few calls are completed successfully, and 90% of initiated calls fail. On average, Random allowed only 8 concurrent calls (Fig. 5). LBT-Simple is a marked improvement – fewer calls are set up, but over 80% of them are completed (Fig. 4), averaging approximately 15 concurrent calls. LBT-DFS initiates more calls, as is expected. However, since the radio conditions at the receiver do not necessarily match those at the transmitter, fewer calls are completed successfully. The average number of concurrent calls is comparable to the simple variant, but still slightly lower.

The local reassignment scheme using our etiquette is midway between the two LBT variants in the number of calls it initiates. However, it has a significantly higher completion rate Fig. 7-d shows that, as expected, all calls setup using our constraints and the local reassignment etiquette adequate C-I. This comes about due to our constraints accurately assessing the tolerable limits of interference temperature in the bandwidth of the calls placed, and preventing any call which would incur or inflict harmful levels of interference from being set up in the first place (the slight dip reported in Fig. 4 is simply because some calls were still in progress when the simulation ended). Comparing Fig. 6-b and Fig. 6-d we see that (except for a dip in Service Ds performance under

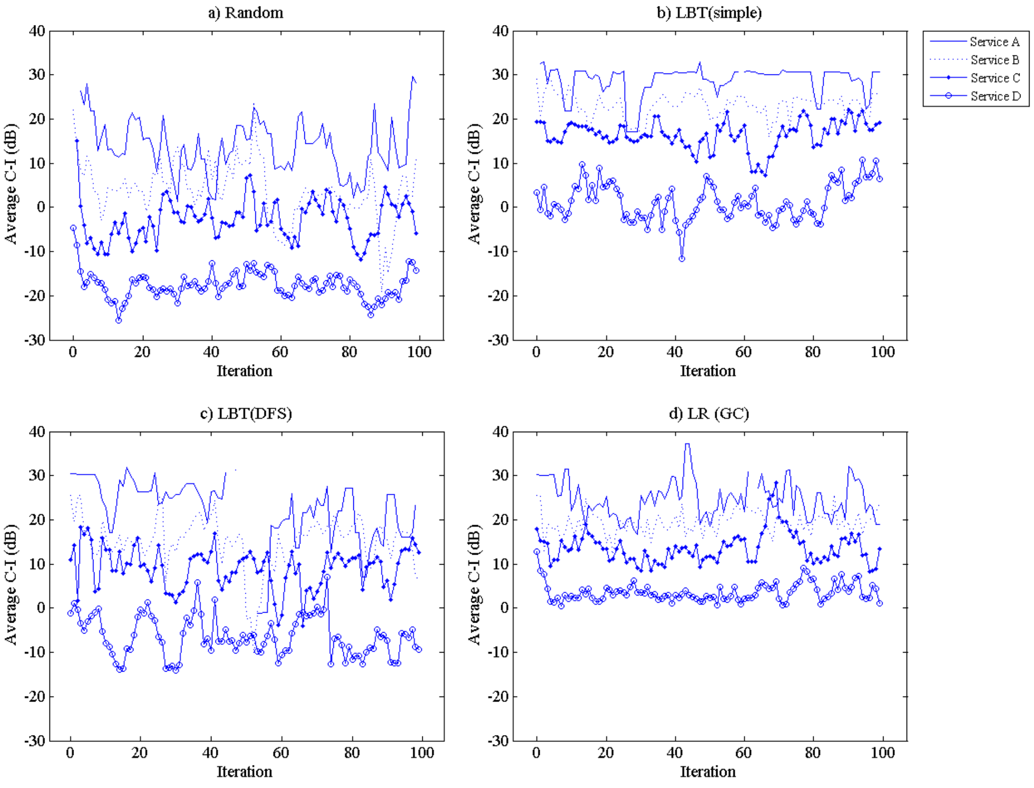

Fig. 6. Average C-I results for different etiquettes (during a typical run)

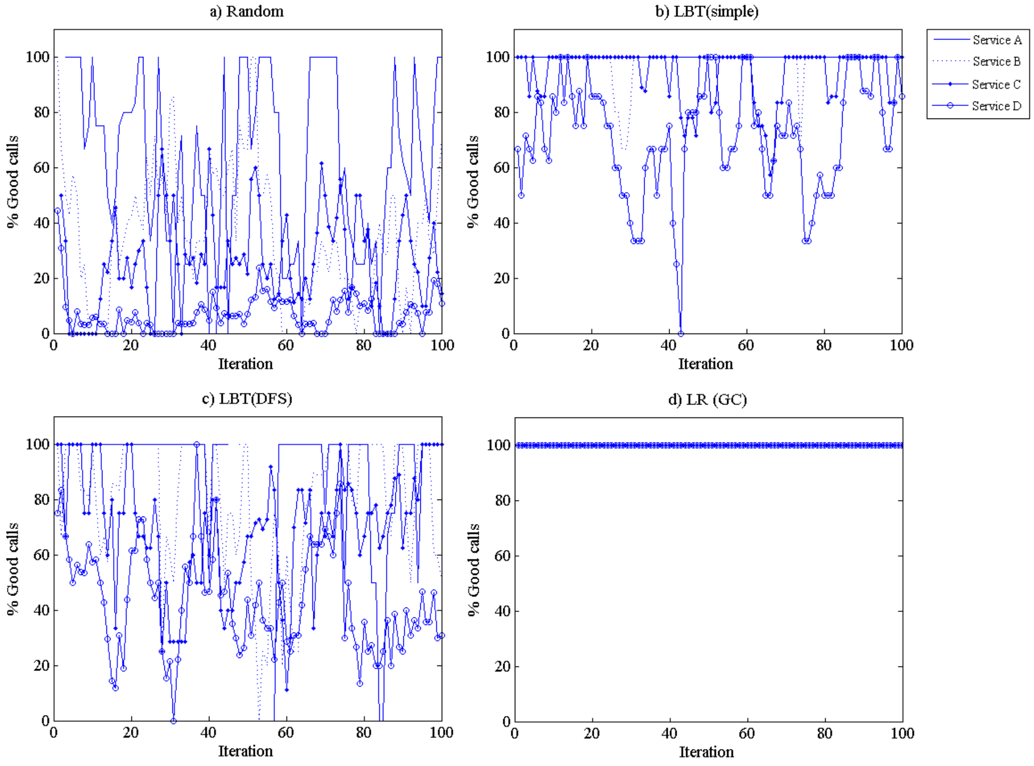

Fig. 7. % admissible quality calls results for different etiquettes (during a typical run)

LBT-simple) that the average C-I achieved is actually lower with our etiquette than that achieved with LBT-simple, again this is due to our constraints accurately assessing the tolerable limits of interference permitted to the calls, leading to more efficient utilisation of the spectrum. On average, our protocol allows over 30 concurrent calls, which is double the number achieved by the LBT variants.

VI. CONCLUSIONS AND FUTURE WORK

Many Cognitive Radio applications search for "holes" in congested spectrum. We have described and evaluated a collaborative spectrum assignment etiquette in which devices not only use, but *create*, spectrum holes in order to facilitate new calls. This is done by dynamically generating specific non-binary link-level interference temperature constraints which faithfully model the potential trade-offs in interference while allowing admissible quality communications, and then using these to re-assign currently transmitting devices. The protocol and constraints are applicable to widely heterogeneous services co-existing in the same spectrum. Our experimental comparison to two common LBT variants has shown that the protocol achieves double the spectrum utilisation in terms of successfully completed calls – although one of the LBT variants initiated more calls, many of the calls suffered unacceptable interference. In all cases, our protocol ensured acceptable call quality throughout the duration of the call.

We have several avenues for future work. The main drawback of our approach is in the time taken to generate and solve the constraints - we are investigating more efficient algorithms, and also the ability to re-use constraints and tuples from previous spectrum reassignments. For example, improved algorithms such as *forward checking with conflict directed backjumping*, developed for binary constraints, could be adapted for the interference temperature constraints used here. Most generally, we intend to consider issues of fairness (see [5]), mobility and the presence of multiple non-cooperating etiquettes.

REFERENCES

[1] M. Cave and W. Webb, "Spectrum licensing and spectrum commons - where to draw the line," in *Wireless Communication Policies and Politics: a Global Perspective*. University of Southern California: Annenberg Research Network, 2004.

[2] J. Peha, "Approaches to Spectrum Sharing," *IEEE Communications Magazine*, pp. 10–11, 2005.

[3] W. Lehr and J. Crowcroft, "Managing Shared Access to a Spectrum Commons," in *IEEE DySPAN'05*, 2005, pp. 134–137.

[4] J. Kruys, "Co-existence of dissimilar wireless systems," 2003, Cisco Systems.

[5] C. Peng, H. Zheng, and B. Zhao, "Utilization and Fairness in Spectrum Assignment for Opportunistic Spectrum Access," *Mobile Network Applications*, 2006.

[6] D. Maldonado, B. Le, A. Hugine, T. Rondeau, and C. Bostian, "Cognitive radio applications to dynamic spectrum allocation: a discussion and an illustrative example," in *IEEE DySPAN'05*, 2005.

[7] D. Raychaudhuri and X. Jing, "A spectrum etiquette protocol for efficient coordination of radio devices in unlicenced bands," in *14th International IEEE PIMRC'03*, 2003.

[8] L. Cao and H. Zheng, "Distributed spectrum allocation via local bargaining," in *IEEE SECON'05*, 2005.

[9] FCC, "In the matter of establishment of an interference temperature metric to quantify and manage interference and to expand available unlicensed operation in certain fixed, mobile and satellite frequency bands," 2003, ET Docket No. 03-237.

[10] P. J. Kolodzy, "Interference temperature: a metric for dynamic spectrum utilization," *International Journal of Network Management*, vol. 16, pp. 103–113, 2006.

[11] T. Clancy, "Formalizing the interference temperature model," *Journal on Wireless Communications and Mobile Computing*, 2007, to appear.

[12] J. Bater, H.-P. Tan, K. Brown, and L. Doyle, "Modelling Interference Temperature Constraints for Spectrum Access in Cognitive Radio Networks," 2007, accepted for *IEEE CogNet Workshop*.

[13] Y. Chen, N. Han, S.-H. Shon, and J. Kim, "Dynamic Frequency Allocation Based on Graph Coloring and Local Bargaining for Multi-Cell WRAN System," in *Proceedings of IEEE Asia-Pacific Conference on Communications (APCC '06)*, 2006.

[14] J. Tamminen, "2.4GHz WLAN radio interface," 2001.

[15] M. Wellens, M. Petrova, J. Riihajärvi, and P. Mähönen, "Building a better wireless mousetrap: Need for more realism in simulations," in *Second Annual Conference on Wireless On-demand Network Systems and Services (WONS'05)*, 2005.

[16] A. Leu, M. McHenry, and B. Mark, "Modeling and analysis of interference in Listen-Before-Talk spectrum access schemes," *International Journal of Network Management*, vol. 16, pp. 131–147, 2006.

[17] D. Grace, A. Burr, and T. Tozer, "A comparison of different distributed channel assignment algorithms for UFDMA," in *Proceedings of the 2nd IEE International Conference on Personal, Mobile, and Spread Spectrum Communications (ICPMSC'96)*, 1996.