Modelling Interference Temperature Constraints for Spectrum Access in Cognitive Radio Networks

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

***Abstract*— With the advent of cognitive radio technology, new paradigms for spectrum access can achieve near-optimal spectrum utilisation by letting each user sense and utilise available spectrum opportunistically while regulating the interference it imposes on other users through interference constraints. However, the simplest and most common forms of such constraints are binary and transmitter-centric, which are often inefficient since they only consider pair-wise sets of transmitters.**

Hence, we propose a non-binary receiver-centric constraint model for spectrum access in cognitive radio networks. Such a model is in line with the recently proposed interference temperature metric that constraints whole subsets of transmitters, thereby permitting interfering signals to be introduced and enabling additional communication, leading to improved spectrum utilisation. These constraints are easy to generate and check, and are currently being used to devise a co-operative negotiated etiquette for cognitive radios offering heterogeneous services in a wireless office networking scenario.

I. Introduction

In today's wireless networks, a *command and control* approach to spectrum management is deployed, where fixed spectrum slices are licensed to each wireless service / technology. However, recent studies [1] have shown that spectrum utilisation is 6.5% (0.8%) and 78% (97%) of spectrum is unutilized in urban (rural) areas. This inefficient use of scarce wireless radio spectrum, along with a dramatic increase in spectrum access for mobile services, have been the driving forces towards new spectrum management paradigms [2].

In the *licensed* model, an exclusive-use license is assigned which may be traded in secondary markets. The licensee is responsible for making all substantive choices as to how the spectrum is used. In contrast, a *non-licensed* model that supports the *coexistence* of *primary* and *secondary* users is enabled by the advent of cognitive radios (CR) [3], [4]. While primary users have priority in spectrum access, secondary users can use available spectrum without interfering with primary users through opportunistic access (overlay) or low power spread-spectrum techniques (underlay). This results in efficient spectrum usage and simplifies deployment of new applications.

However, the requirement for spectrum efficiency certainly predates the advent of CR, e.g. [5]. There is a wealth of literature on solving the Frequency Assignment Problem (FAP) in cellular networks, key to which is in modelling this problem as generalised Graph Colouring (GC). The constraints used in FAP studies were largely *binary* (restricting the assignments of frequency on pairs of transmitters), and were usually derived from a *re-use distance*, or an estimation of the effect of interference on the cell's receivers from another potential interferer [6]. In [7], non-binary constraints were constructed which considered the effects of multiple sources of interference from across the network; these constraints placed restrictions on the simultaneous spectrum assignment for an arbitrary number of transmitters to ensure that the receivers in the cell all maintained communications of adequate quality.

In this paper, we refine these constraints and develop a framework for constraint-based approaches to spectrum access in CR networks. In particular, we propose a non-binary receiver-centric constraint model in line with the recently proposed interference temperature metric. We demonstrate that our proposed approach may improve spectrum utilisation compared to traditional transmitter-centric approaches in a non-licensed spectrum regime.

II. A Framework for Constraint-based Approaches to Spectrum Access

We consider a distributed deployment of CR-enabled secondary users over a geographical region, self-configured to form an ad-hoc network and access (or share) spectrum in a *co-operative* manner. Without loss of generality, we assume that the spectrum band is divided into discrete non-overlapping and non-interfering channels in the frequency domain.

Each user has the ability to sense the radio environment and determine the available spectrum. As with most proposed spectrum access schemes, we assume that each user has access to a perfectly synchronized and dedicated (interference-free) Common Signalling Control Channel (CSCC) from which parameters required for spectrum access can be computed through exchange of signalling messages. Although this is a restrictive assumption, it is commonly adopted in the literature of dynamic spectrum access.

Each user then *reconfigures* itself, e.g., in terms of transmission power [8], channel [9] or a combination of both [10] to maximize the spectrum utilisation while regulating the interference it imposes on other users. We consider a *co-channel* interference model, where adjacent channel signals are entirely rejected at each receiver.

Interference is typically and traditionally regulated in a *transmitter-centric* way, which means interference can be controlled at the transmitter through the transmitted power, the out-of-band emissions and location of individual transmitters. Based on each user's transmission power as well as the topology of the network, *interference constraints* are constructed and used to determine the spectrum *assignment* to each user such that interference remains within acceptable levels. In this section, we formalise the framework for spectrum access using constraint-based models.

A. Terminology

A *constraint* consists of a *scope*, S, which is a subset of the variables in a problem; and a *relation*, R, which is a function or expression describing the simultaneously allowed (or disallowed) assignments of values to variables in the scope:

$$C = \langle S, R \rangle.$$

The relation can be expressed extensionally (i.e. as sets of values), or intensionally (i.e. a formulaic expression). In the context of spectrum access, the scope for a user is the set of transmitting users in its vicinity that may potentially interfere with its ongoing communication with another user; and the relation may specify, say, the channel separations required between the user and its interferers to maintain interference to within acceptable levels.

Constraints are generally described in terms of their *arity*, the number of variables in their scope. *Binary* constraints place restrictions on the simultaneously assign-able values to particular *pairs* of variables. A binary constraint problem is modelled using *only* binary constraints. Binary constraint models are most common in the literature ([11] and [12] are good starting treatises on the subject), but there is increasing interest in *non-binary* constraints which tackle larger subsets of variables in a particular problem than just two [13].

B. Transmitter-centric Constraints

The simplest and most common form of interference constraints are binary and transmitter-centric, and they are used widely in the literature on spectrum access in CR networks [14], [9]. In essence, these constraints define a *re-use distance* between any *pair* of transmitters; within this distance, the re-use of a channel, or set of channels is not permitted.

Using the notations in [9], we define $d_s(t_i,c)$ as the *detection range* (for a receiver) of transmitter t_i in channel c, where $d_s(t_i,c)$ increases with the transmission power of user t_i and is user- and channel-dependent in general since it is constrained by the activity of primary users.

For any transmitter pair (t_i,t_j), the re-use distance for channel c is then given by $d_s(t_i,c)$ + $d_s(t_j,c)$. Accordingly, if $Dist(t_i,t_j)$ is the distance between t_i and t_j, they can share (or re-use) channel c only if the following condition holds:

$$Dist(t_i, t_j) > d_s(t_i, c) + d_s(t_j, c). \quad (1)$$

The above constraint eliminates the possibility of potential interference to receiver r_i (r_j) from t_j (t_i), where r_i (r_j) is the intended receiver and within the detection range of t_i (t_j). To illustrate, let us consider a network with 3 transmitting users (nodes), $\{t_1, t_2, t_3\}$ sharing 3 channels, $\{A, B, C\}$ as shown in Fig. 1(a), where the detection range of each transmitter is given by the radius of the dotted circle around it. According to Eq. (1), transmitters t_2 and t_1 cannot use channel C simultaneously while t_1 and t_3 can.

By mapping each channel into a colour, binary transmitter-centric constraints such as Eq. (1) can be abstracted into a graph colouring (GC) model [9], based on which channels (colours) can be assigned to transmitters in a CR network. The corresponding GC model for the scenario in Fig. 1(a) is shown in Fig. 1(b). A label on edge $t_i - t_j$ indicates channel(s) unusable simultaneously by transmitters t_i and t_j according to Eq. (1). Accordingly, a feasible assignment is given by $\{(t_1$, *A*), (t_2, *B*), (t_3, *C*)$\}$, where (t_i, *J*) indicates that node t_i is assigned channel *J*.

C. Receiver-centric Constraints

Although interference constraints for spectrum assignment are typically constructed in a transmitter-centric way to alleviate *co-channel* interference, interference *actually* takes place at the *receivers*. Therefore, recently, a new metric for measuring interference at the receiver, known as *interference temperature* has been suggested by the FCC [15]. While there is still controversy over its feasibility and usefulness, we attempt to demonstrate its merits towards a constraint-based approach for dynamic spectrum access.

To illustrate the interference temperature metric, we consider the scenario in Fig. 1(a) and assume that user t_1 (t_2) is transmitting in channel *A* (*B*) to r_1 (r_2) with received power P_1 (P_2) dB over a noise floor NF_A (NF_B). The quality of the transmission is usually quantified in terms of the carrier-to-interference ($C\text{-}I_{i,c}$) ratio, which measures the ratio of the *desired* received power, P_i from transmitter t_i at r_i in channel c to the sum of *unwanted* (or *interfering*) co-channel received signal power, I_c, and the noise floor, NF_c. We define the communication between t_i and r_i in channel c to be *admissible* if $C\text{-}I_{i,c} > C\text{-}I_{th,i,c}$ (termed C-I *threshold*), or equivalently, $I_c < P_i$ - $C\text{-}I_{th,i,c}$ - NF_c, i.e., the allowable interference in channel c is given by P_i - $C\text{-}I_{th,i,c}$ - NF_c. This is illustrated in Fig. 2 for the scenario in Fig. 1(a).

By allowing additional interference at each receiver, additional communication links can be supported in its vicinity for a given spectrum availability, giving rise to improved spectrum utilisation. For example, let's assume that user t_1 is transmitting to r_1 in channel *A* according to the scenario in Fig. 1(a). Using binary transmitter-centric constraints, according to Fig. 1(b), user t_3 will be prohibited from transmitting to r_3 in channel *A*. However, with the receiver-centric interference temperature constraint, as long as the co-channel interference associated with user t_3's transmission is less than P_1 - $C\text{-}I_{th,1,A}$ - NF_A, user t_3 will be able to transmit to r_3 in channel *A* as well.

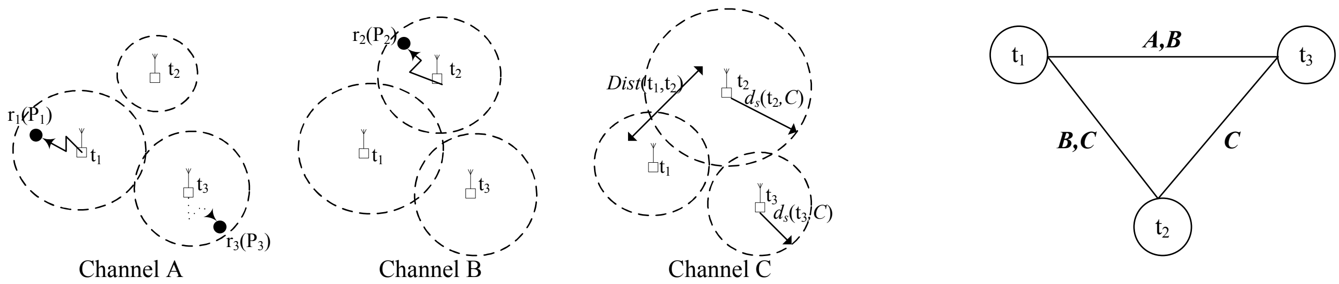

(a) Transmitter-centric interference constraints

(b) Colour-sensitive Graph Coloring Model

Fig. 1. (a) An illustration of binary and transmitter-centric interference constraints and (b) the corresponding colour-sensitive graph colouring model for allocating 3 channels, $\{A, B, C\}$ amongst 3 transmitting users, $\{t_1, t_2, t_3\}$ (represented by vertices). Each dotted circle represents the interference range of a node and the label on edge *i*—*j* indicates spectrum unusable by nodes *i* and *j* simultaneously.

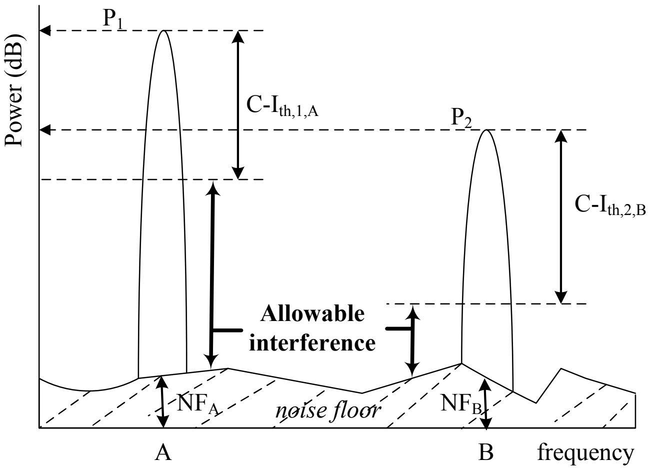

Fig. 2. An interference temperature model to determine the additional allowable co-channel interference to enable additional communication without degrading ongoing communication. This is given by P_i - *C-*$I_{th,i,c}$ - *NF*$_c$, where P_i is the received power of the ongoing transmission from user *i* in channel *c*, *NF*$_c$ is the noise floor and *C-*$I_{th,i,c}$ is the carrier-to-interference ratio threshold for admissible call quality.

1) Binary Receiver-centric Constraints: Having illustrated the interference temperature metric, we now have a number of options as to exactly how to map this into receiver-centric interference constraints. For t_i transmitting in channel c with received power P_i at receiver r_i, the maximum tolerable amount of co-channel interference at r_i is P_i - *C-*$I_{th,i,c}$ - *NF*$_c$. Accordingly, we can generate a binary constraint (still akin to GC models) such that no transmitter t_j subjects r_i to interference beyond P_i - *C-*$I_{th,i,c}$ - *NF*$_c$. In other words, t_j can share channel c with t_i only if the following condition holds:

$$P_{j,i} < P_i - C\text{-}I_{th,i,c} - NF_c, \tag{2}$$

where $P_{j,i}$ is the received interference power at r_i from t_j. This constraint can be re-written in terms of an *exclusion* distance (similar to the *re-use* distance for binary transmitter-centric constraints) as follows:

$$Dist(t_j, r_i) > f(P_i - C\text{-}I_{th,i,c} - NF_c),$$

where $f()$ is a *decreasing* function. In contrast to the re-use distance, the exclusion distance is based on a more realistic relationship of the strength of the wanted signal, the C-I threshold (giving us the tolerable interference), and whether a potential interferer would provide too much interference.

Consider a multiple-interferer scenario depicted in Fig. 3, where r_1 is receiving a signal from t_1 while being (co-channel) interfered by three other transmitters, t_2, t_3 and t_4. Using Eq. (1), the scope of each binary transmitter-centric constraint is depicted by the respective dotted circle in Fig. 3(a). The corresponding scope of each binary receiver-centric constraint, according to Eq. (2), is shown in Fig. 3(b). Assuming symmetric transmissions from the interferers and Eq. (2) is satisfied, we have the following situations:

1) If each interferer contributes up to one third of the total tolerable interference at r_1, the communication between t_1 and r_1 will remain in good quality;
2) If each interferer contributes up to half of the total tolerable interference at r_1, then only *two* of the interferers may continue to interfere for good quality communication between t_1 and r_1 to endure;
3) If each interferer contributes more than half of the total tolerable interference at r_1, then only *one* of the interferers may continue to interfere for good quality communication between t_1 and r_1 to endure.

Hence, we observe that binary receiver-centric constraints may permit channel assignments that result in inadequate C-I for communication. However, we might still place binary constraints that guarantee the total received interference to be below the interference temperature limit, provided we know in advance which pair(s) to restrict. But, in doing so, we may possibly discard solutions to the local problem which may in fact be desirable.

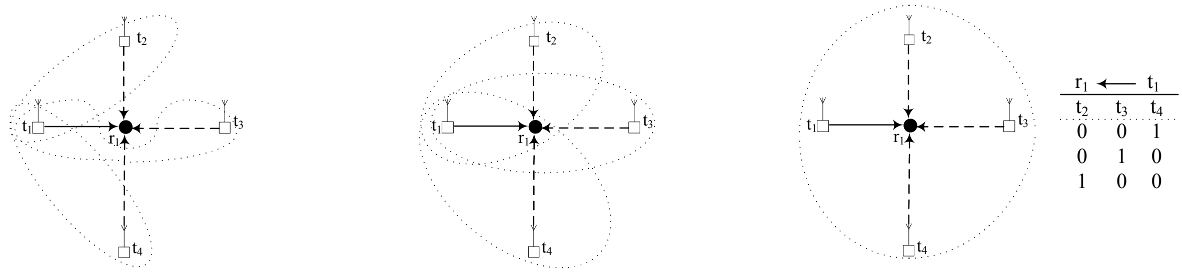

(a) binary, transmitter-centric (b) binary, receiver-centric (c) non-binary, receiver-centric

Fig. 3. Comparison of interference constraints for a multiple-interferer scenario: (a) Scope of binary transmitter-centric constraints (b) Scope of binary receiver-centric constraints and (c) Scope of non-binary receiver-centric constraints and an example of a relation comprising three tuples (right).

2) Non-binary Receiver-centric Constraints: To overcome the shortcomings of binary constraints, we can instead surround several transmitters with a *hyperedge*, and form a non-binary receiver-centric constraint as shown in Fig. 3(c). While such constraints have been considered previously for cellular network problems [7], we refine them here for use in CR networks.

Let us assume that each interferer provides slightly over one third of the receiver r_1's tolerable interference (which implies that binary constraints will lead to excessive interference at r_1 according to Section II-C.1). Here the potential interferers t_2, t_3 and t_4 are permitted to be co-channel (have a channel separation of 0 channels), or non-interfering (have a channel separation of at least 1 channel) with t_1.

We can now generate a constraint consisting of a set of tuples each specifying an interference relation on the scope of our constraint (here t_2, t_3 and t_4). Each tuple expresses the minimum separation between channel sets of the wanted signal, and that of the unwanted signal from each interferer, necessary to maintain required C-I. (The interference budget can potentially be "filled" in a number of ways - actually it doesn't even need to be filled). An example constraint with a relation that comprises three tuples is shown in Fig. 3(c).

Let us consider the first tuple (t_2:0, t_3:0, t_4:1). The minimum separation at t_2 and t_3 is ≥ 0 (i.e. a co-channel assignment is permitted somewhere on their channels) *if and only if* t_4 has a separation of ≥ 1 (i.e. non-interfering) with t_1. This simple example is symmetrical and so two of the three devices may interfere *provided* that the remaining one does not. Should any of the transmitters have larger separations in channel space (especially those permitted to be co-channel) then they contribute less interference so the C-I of the received signal at r_1 increases.

Due to space constraints, the method for generating constraints will not be described in detail here; however, it should be noted that a search similar to Depth-first search with an enhanced backtracking step designed to explicitly avoid searching redundant areas of the search tree is used.

III. Comparison of Models

In this section, we demonstrate the merits of non-binary receiver-centric constraint models for spectrum access, and use this to further justify our position. We will ignore the simple binary receiver-centric constraints of Section II-C.1. Though they may work in practice if there is a further requirement that channel use is well spread out throughout the available spectrum band, solutions to these constraints are not necessarily suitable assignments in a network.

We term the binary transmitter-centric constraints as *"Cautious"* (or CT), since it eliminates the possibility of interference from other transmitters. We term the non-binary receiver-centric constraints as *"Interference Temperature"* (or IT), since it accurately maps the possible interference contributions at a link level to ensure that the interference temperature for that link is not met or exceeded.

For a given CR network that comprises N transmitter/receiver pairs $\{t_i, r_i\}_{i=1}^{N}$, let $\mathbf{A}_C = \{a_1, a_2, \cdots, a_N\}$ be the set of spectrum assignments that satisfy the interference constraints C. We consider the following questions:

1) *Does any spectrum assignment* $a_{\mathrm{CT}} \in \boldsymbol{A}_{\mathrm{CT}} \implies a_{\mathrm{CT}} \in \boldsymbol{A}_{\mathrm{IT}}$*?*
The answer is clearly yes. CT does not permit interference from any other user at all, whereas IT does, provided the total interference is below that tolerable by the receiver (for a given signal level and C-I threshold). In the worst case when IT is as restrictive as CT, the set of potential interferers = $\{\ \}$. Therefore, in the C-I equation, the interference term is zero, hence C-I is ∞ which is greater than any likely C-I threshold.
2) *Does any spectrum assignment* $a_{\mathrm{IT}} \in \boldsymbol{A}_{\mathrm{IT}}$ *exist such that* $a_{\mathrm{IT}} \notin \boldsymbol{A}_{\mathrm{CT}}$*?*
Again the answer is clearly yes. Suppose we have a receiver r_j whose received wanted signal power is P_j. For a particular desired C-I threshold = θ, we have a spectrum assignment $a_{\mathrm{IT}} \notin A_{\mathrm{CT}}$ as long as the sum of unwanted interference powers, I_j is such that $0 < I_j < \frac{P_j}{\theta}$.

We can infer from the above discussion that spectrum access using IT constraints achieves *better* (or at worst, equal) spectrum utilisation than (as) that using the CT constraints.

To support the above inference numerically, we simulate a *homogeneous* CR network that comprises N={10,20,$\cdots$,50} transmitters deployed randomly over a geographical area of 100m $\times$ 100m. Each transmitter t_i has a common detection range, d_s, leading to a common re-use distance, $2d_s$; the corresponding receiver r_i is deployed randomly such that *Dist*(t_i,r_i)$\leq d_s$. Each communication pair has a common C-I threshold, θ = {3,6,$\cdots$,15} dB and the transmission power is the minimum level required to satisfy the following condition:

$$P_{min} = NF + 20dB,$$

where the 20dB margin ensures that the C-I threshold will *always* be met while possibly allowing additional interference. Next, we generate the interference constraints C = {CT, IT} and construct the channel assignment $m_{C,N,\theta}$, from which we can determine the minimum number of channels needed, $\mathit{NumCh}_{C,N,\theta}$, to solve C.

We plot $\mathit{NumCh}_{C,N,\theta}$ (averaged over 1000 simulation runs) as a function of N for θ = 9 dB and as a function of θ for N = 30 in Fig. 4. We observe that, for a given (N, θ), IT result in a lower spectrum requirement than CT since IT permit some interference (subject to the C-I threshold) while CT do not permit any co-channel interference at all.

As the network becomes more congested (i.e., as N increases for fixed θ), the spectrum requirement increases using both types of constraints. This is expected since the likelihood of mutual interference increases as more users are communicating within a fixed geographical area. Since IT permit some interference (subject to the C-I threshold) while CT do not permit any co-channel interference at all, the increase in spectrum requirement is more gradual with IT and the gain achieved with IT compared with CT increases as the network becomes more congested.

However, as the requirement for admissible quality communication becomes more stringent (i.e., as θ increases for fixed N), the spectrum requirement using CT remains invariant since these constraints depend only on d_s, and are independent of θ. On the other hand, the spectrum requirement using IT increases with θ since the level of permissible co-channel interference decreases and hence, channel re-use is reduced.

To better quantify spectrum efficiency, we define the following metric, $\eta_{C,N,\theta}$:

$$\eta_{C,N,\theta} = \frac{N}{\mathit{NumCh}_{C,N,\theta}},$$

which evaluates the number of communication pairs that can be supported by each channel (similar to the definition used in [16]). Hence, the *higher* the value of $\eta_{C,N,\theta}$, the *better* the spectrum efficiency that can be achieved with constraints C. We plot $\eta_{C,N,\theta}$ as a function of N and θ in Fig. 5.

While the spectrum efficiency achieved by using CT remains almost invariant ($\approx$ 2) with (N, θ), with IT, the spectrum efficiency improves as the network becomes congested (N increases) and the requirement for acceptable quality communications becomes relaxed (θ decreases). However, we expect the spectrum efficiency to be degraded for both approaches once N increases beyond a certain threshold, N_{thres}. In addition, since the receivers are involved in generating receiver-centric constraints, we expect higher levels of overhead with the IT approach due to exchange of signalling messages over the CSCC.

More extensive numerical results covering broader network scenarios as well as a rigourous proof for the over-cautiousness of the CT constraints can be found in the full version of the paper [17].

IV. Conclusions and Future Work

With the advent of cognitive radio technology, new paradigms for spectrum access can achieve near-optimal spectrum utilisation by letting each user sense and utilise available spectrum opportunistically while regulating the interference it imposes on other users. Although the simplest and most common form of interference constraints are binary and transmitter-centric, they are often inefficient since they only consider pair-wise sets of transmitters. Hence, we propose a non-binary receiver-centric constraint model for spectrum access in cognitive radio networks. Such a model is in line with the recently proposed interference temperature metric that constraints whole subsets of transmitters, thereby permitting interfering signals to be introduced and enabling additional communication, leading to improved spectrum utilisation. This is shown in this paper through numerical experiments. However, since the receivers are involved in generating receiver-centric constraints, we expect higher levels of overhead compared to transmitter-centric approaches.

We have implemented an efficient model of non-binary constraints which are simple and fast to both generate and test. Our code is able to generate constraints of arbitrary arity (i.e. for random networks), and whose C-I relation is constructed for different thresholds (i.e. supporting heterogeneous users with different service requirements).

At present, we are using these constraints to devise a negotiated etiquette for CRs offering heterogeneous services in a wireless office scenario. While this extends and builds on the work in [18], the following differences are noted:

1) As opposed to binary transmitter-centric constraints, we apply non-binary receiver-centric constraints that more accurately models the interference, resulting in better spectrum utilisation;
2) Instead of assuming continuously backlogged homogeneous users, we consider a highly dynamic traffic model for users that operate over a number of dissimilar services (each requiring different minimum C-I levels and having different effective bandwidths).

References

[1] A. Petrin and P. G. Steffes, “Analysis and comparison of spectrum measurements performed in urban and rural areas to determine the total amount of spectrum usage,” *Proc. of the ISART*, March 2005.

[2] W. Lehr and J. Crowcroft, “Managing shared access to a spectrum commons,” *Proc. of the IEEE DySPAN*, pp. 420–444, November 2005.

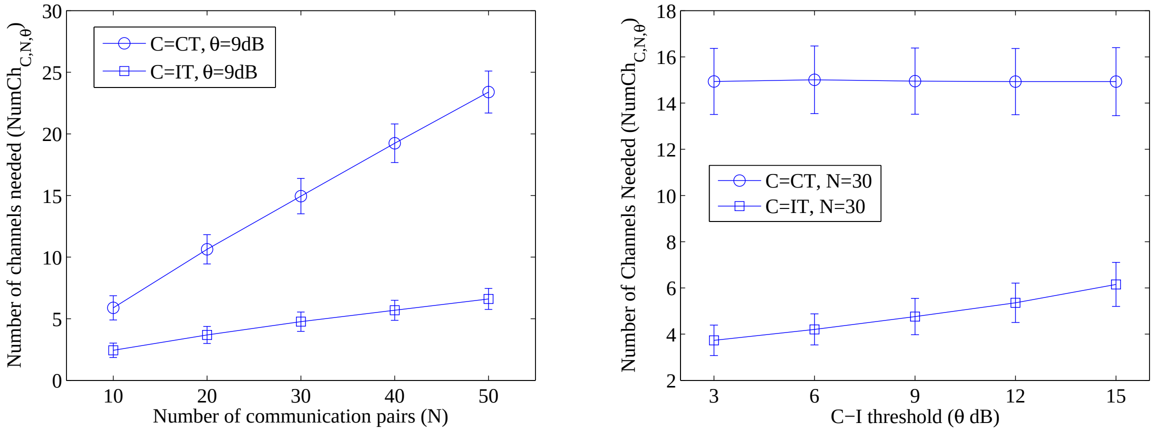

Fig. 4. Comparison of the number of channels needed ($NumCh_{C,N,\theta}$) to ensure acceptable communication (C-I threshold θ dB) amongst N communication pairs using CauTious binary transmitter-centric constraints (CT) and Interference Temperature based non-binary receiver-centric constraints (IT) for channel assignment in a CR network.

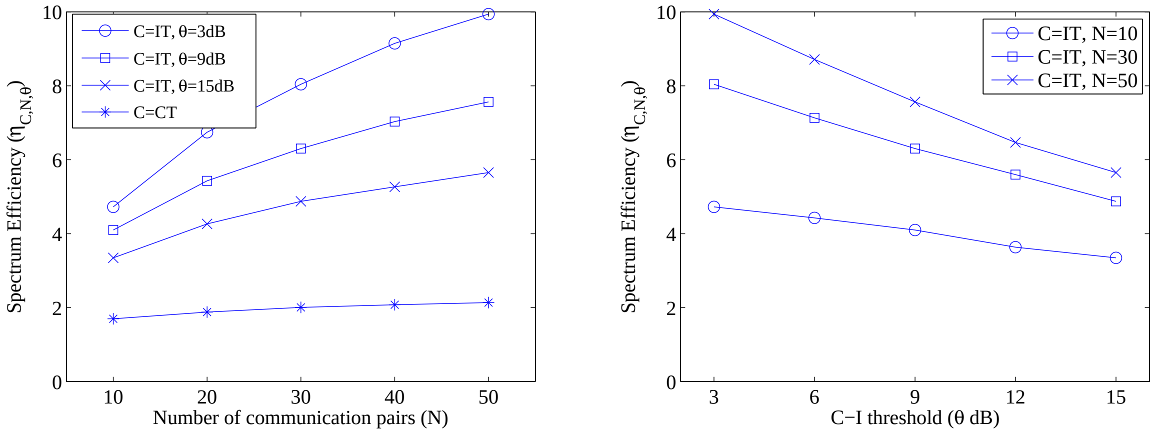

Fig. 5. Comparison of the spectrum efficiency ($\eta_{C,N,\theta}$ = number of communication pairs where acceptable communication can be sustained (C-I threshold = θ dB) per channel) using CauTious binary transmitter-centric constraints (CT) and Interference Temperature based non-binary receiver-centric constraints (IT) for channel assignment in a CR network with N communication pairs.

[3] J. M. III, "Cognitive Radio: An Integrated Agent Architecture for Software Defined Radio," Ph.D. dissertation, KTH Royal Institute of Technology, 2000.

[4] S. Haykin, "Cognitive Radio: Brain-Empowered Wireless Communications," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 2, pp. 201–220, February 2005.

[5] L. Berry, "Spectrum metrics and spectrum efficiency," *IEEE Trans. Electromagnetic Compatability*, vol. EMC19, no. 3, pp. 254–260, 1977.

[6] W. Hale, "Frequency assignment: Theory and applications," *Proc. of the IEEE*, vol. 68, pp. 1497–1514, 1980.

[7] J. Bater, "Non-binary (higher-order) modeling and solution techniques for frequency assignment in mobile communication networks," Ph.D. dissertation, University of London, 2000.

[8] J. Neel, J. Reed, and A. MacKenzie, "Cognitive radio network performance analysis," *Cognitive Radio Technology*, August 2006.

[9] C. Peng, H. Zheng, and B. Zhao, "Utilization and fairness in spectrum assignment for opportunistic spectrum access," *Mobile Networks and Applications*, vol. 11, no. 4, pp. 555–576, August 2006.

[10] N. C. Ericsson, S. Falahati, A. Ahlen, and A. Svensson, "Scheduling and Adaptive Transmission for the Downlink in 4G Systems," *Proc. of the IEEE DySPAN*, pp. 251–258, November 2005.

[11] E. P. K. Tsang, *Foundations of Constraint Satisfaction*. Academic Press, 1993.

[12] K. Marriot and P. J. Stuckey, *Programming with Constraints*. MIT Press, 1999.

[13] J. C. Regin, *Constraints and Integer Programming Combined*. Kluwer, 2003, ch. Global Constraints and Filtering Algorithms.

[14] H. Zheng and C. Peng, "Collaboration and fairness in opportunistic spectrum access," *Proc. of the IEEE ICC*, vol. 5, pp. 3132–3136, May 2005.

[15] P. J. Kolodzy, "Interference temperature: a metric for dynamic spectrum utilization," *International Journal of Network Management*, vol. 16, no. 2, pp. 103–113, March 2006.

[16] F. Berggren, "Capacity analysis for non-cooperative interference environments," *Proc. of IEEE VTC*, vol. 3, pp. 1479–1483, June 2005.

[17] J. Bater, H. P. Tan, K. Brown, C. O'Sullivan, and L. Doyle, "An interference temperature constraints model for spectrum access in cognitive radio networks," *Submitted to Special Issue of IEEE JSAC in Cognitive Radio: Theory and Applications*, March 2007.

[18] L. Cao and H. Zheng, "Distributed spectrum allocation via local bargaining," *Proc. of the IEEE SECON*, pp. 475–486, September 2005.