

On the Impact of Introducing Advanced Devices into a Cognitive Radio Network

Joe Bater*, Kenneth N Brown* and Linda Doyle†

Centre for Telecommunications Value-chain Research

*Cork Constraint Computation Centre, University College Cork, Republic of Ireland

Email: j.bater@4c.ucc.ie, k.brown@4c.ucc.ie

†Dept. Electronic Engineering, Trinity College Dublin, Republic of Ireland

Email: ledoyle@tcd.ie

Abstract—Cognitive Radio promises better spectrum utilisation through decentralised control, sensing and decision making. The devices are intended to be inexpensive, upgradeable and adaptable. However, introducing upgraded devices into a decentralised environment may cause unexpected results, the upgrades may require a critical mass to show any benefit, and may have an adverse impact on existing devices, which may deter their adoption. We study, in simulation, the gradual introduction of cognitive devices into an existing population using a spectrum access etiquette, in three different scenarios for the device capability. We show that the new devices gain immediate benefits, and that there is always an incentive to upgrade. We show that overall spectrum efficiency is improved. Finally, we show that there is no negative impact on existing devices, and that the introduction of cognitive devices may even improve the success rate of lesser devices.

I. INTRODUCTION

Cognitive Radio [1] has received increasing attention in recent years, as part of the drive toward better spectrum management. Traditionally, blocks of spectrum have been allocated for specified services and specified technologies. In latter years dynamic spectrum access, advanced common regimes based on sophisticated spectrum etiquettes and exclusive usage rights scenarios that support technology and service neutral regimes have challenged the traditional approach. A cognitive radio is one which can sense its environment and adapt its operating parameters accordingly, and is seen as an enabling technology in all of these cases. But introducing cognitive radio into existing communication regimes may prove difficult [2]: the value of cognitive radio must be demonstrated and users must have an incentive to switch, but there are regulatory concerns that the introduction should not have a significant negative impact on existing applications and services or on users of existing devices.

In this paper, we begin to study the impact of introducing cognitive radio. We consider three questions:

- 1) can we quantify the advantage of cognitive devices for individual users to promote early adoption?
- 2) can we identify the general social benefit of introducing cognitive radios?
- 3) what is the impact on existing devices as cognitive radios become more prevalent?

To do this, we simulate the gradual migration of users in a population from existing legacy devices to radios with increased capability, in three scenarios with different levels of capability, and measure the impact. We show that there is an immediate incentive for users to switch to cognitive radio - in all cases, at all stages, the cognitive radio has a higher success rate than the legacy devices, even when there are few other cognitive radios in the population. We show that there is an immediate social benefit - in all cases, the total spectrum efficiency rises and continues to rise as we introduce more cognitive radios. Finally, we show that there appears to be no negative impact on the legacy devices - although their share of spectrum drops, the success rate for individual legacy devices actually rises, as the cognitive devices are able to create space to allow more of their calls.

II. BACKGROUND

Mitola coined the term *Cognitive Radio* [1] to describe devices which were *aware* of their radio environment and could *decide* (on the basis of these observations) to *act* accordingly without intervention or manipulation by the user. Such devices are able to operate across widely ranging frequency bands and employ a variety of modulation and coding schemes. Further suggestions to relieve pressure on heavily congested spectrum set aside for telecommunications use are to allow access to under-utilised spectrum notionally belonging to other services. This latter suggestion has caused a great deal of interest in the CR community, in which “secondary” unlicensed users might opportunistically access the spectrum of a “primary” or licensed user, provide that they vacate the spectrum in a timely manner should the primary user require it.

Whether in the guise of *multi-mode communications* in which SDR devices are able to switch between programs (algorithms) implementing multiple radio-access-technology (multi-RAT) implemented in a single hardware e.g. [3], [4], [5]; *transmission power control* has been shown to enable devices to reduce potential interference on other devices, if higher transmit powers are not necessary then devices should step their transmit powers down accordingly [4], [6]; re-configuration of the *modulation scheme* adaptive to the user requirements and channel conditions is another feature suggested for the cognitive capability of CRs [6] – in varying

radio conditions it might be better to consider trading achieved data-rate and bit-error-rate (BER) by altering the transmissions modulation scheme; finally *frequency agility* has underpinned dynamic spectrum access in which transmissions are able to jump to unused/better operating frequencies e.g. “white space”. This has been the fundamental requirement for negotiation etiquettes in which network devices are able to alter their assignments to maximise spectrum utility e.g. [7], [8], [9].

The purpose of a spectrum etiquette is, in the absence of a formal protocol, to describe a manner of “polite” behaviour among devices in a network. These might range from the Part.15 rules [10] which specify technical requirements and limits on the broadcast power of devices using unlicensed spectrum, to *listen-before-talk* protocols in order to reduce CSMA collisions e.g. [11]. More advanced etiquettes can request that devices change their own spectrum assignments based on the requests of other devices e.g. [7], [8], [9] - among the key issues here are fairness in the resulting assignment and maximising spectrum utilisation.

The most widespread model for spectrum assignment is based on an analogy to the GRAPH COLOURING problem arising in Discrete Mathematics. Here nodes in a graph represent devices and colours different frequency channels. Since radio signal power attenuates rapidly with distance the likelihood of a device interfering with another depends on proximity of both frequency used and geographical location. Consequently edges are drawn on the graph between nearby devices and co-channel frequency re-use between these devices is prohibited.

The most common constraints in spectrum utilisation literature are *binary* since they are specified in terms of a *pair* of devices. If any given two devices are placed within a defined *re-use distance*, they should not be permitted to simultaneously transmit on the same channel. Thus each device considers its immediate neighbours individually to construct an *interference-graph* or *conflict-graph*, G . Here transmitting devices are represented as a set of vertices, V , and the edge set $e(v_i, v_j) \in E$ represents an interference constraint such that v_i and v_j may not have co-channel assignments (due to the interference on the receiver/s served). The direct analogy between channel assignment and GRAPH COLOURING is widely credited to Metzger [12] (e.g. [13], [14]), and this approach remains the dominant model today, e.g. [7], [15], [16]. Many methods are possible to determine re-use distances - one of the most recent is due to Cheng *et al.* [13], and this will be discussed in Section III. Heuristics for GC problems have been widely adapted to channel assignment, e.g. [7], [15], [16].

In practice, the quality of a communications link is determined by the *bit error rate* (BER), which can be predicted by the ratio of the received wanted signal strength to that of the sum of the interference suffered (based on the modulation type and amount of error-coding). The *signal to interference (plus noise) ratio* is generally given as:

$$\text{SINR} = \frac{p_t G_{tr}}{\sum_{i, i \neq t}^N p_i G_{ir} I(t, i) + N}$$

where p_t is the transmit power of transmitter t ; G_{tr} is the attenuation loss between t and the receiver r ; $I(t, i)$ is 1 if t and the interferer i are co-channel and 0 otherwise; N is the background noise. In [17], it was shown that binary GRAPH-COLOURING constraints can be inefficient in terms of spectrum usage - that is, they deny assignments that would give acceptable call quality. Instead, a constraint model was proposed to consider and strictly bound the maximum allowed interference within the bandwidth of the call (rather than as a general measure of “interference” across a wide spectrum band).

It is likely that many different types of traffic, with widely different bandwidth, power and range requirements, will have to co-exist in future shared radio bands [18], [19]. While there is increasing interest in etiquettes which allow dissimilar or heterogeneous services / multiple radio access technologies (multi-RAT), e.g. 802.11a/b/g, 802.16, Bluetooth and UWB ([20]) sharing the spectrum with equal priority and status, most proposed (and implemented) etiquettes only consider homogeneous sets of devices/services.

III. THE BASE LEVEL RADIOS

We begin by defining the radio environment and the devices that will form our legacy population, similar to devices studied in [7] and [9]. In fact, these devices already have significant cognitive capability. It should be noted that this approach still represents a state-of-the-art in resource allocation problems for cognitive networks.

A. Heterogeneous Service Model

We assume devices are able to select from multiple technologies at the start of a call (and they then stay with service for its duration), and so we define a representative set of three different abstract services, shown in I. The multi-channel services require a contiguous block of channels. The individual channels may be altered dynamically during the call.

TABLE I
SPECIFICATION OF REQUIREMENTS FOR DISSIMILAR SERVICES

Service	No. of Channels	Target SINR (dB)
A	1	12
B	2	9
C	5	6

B. Power control

We assume devices transmit at a fixed power (akin to a WLAN or GSM system [21]). Here the fixed transmit power is set to 10dBm.

C. Frequency agility

We assume devices have a fixed frequency range of 10 channels.

D. Interference Model

In order to determine what channels can be used, each device needs a model of acceptable interference. We assume a model based on simple graphs, in which constraints forbidding co-channel assignments are placed between pairs of graphs, depending on their distance apart, their SINR requirements, and their transmit powers. We use the *interFERENCE Double Disc* (FDD) model, recently proposed by Cheng et al. [13]. FDD graphs place two concentric rings about a transmitting device with the first (inner) disc enclosing a radius, r_v , in which the signal strength of the transmitted signal is viable, and the second (outer) ring encapsulating an area in which the signal is too weak to use, but strong enough to act as an interference to other receivers, R_v . FDD graphs place edges between pairs of vertices *iff* R_v of one device overlaps with r_v on another.

In Fig. 1 the devices are operating on different transmit powers, and in order to ensure that device y 's interference region does not overlap with the device x 's transmission region, an edge $e(v_x, v_y)$ is now added to E *iff*:

$$\text{Dist}(x, y) \leq \text{Max}(R_x + r_y, r_x, R_y)$$




Fig. 1. FDD Graphs for Devices with Different Transmit Powers

E. Graph Colouring Algorithms

Once each device constructs a model of the environment, it needs to determine which channels can be assigned to each device, satisfying the graph constraints. We assume a standard approach (shown in Algorithm III-E), based on an adaptation of the Brelaz colouring heuristic [22]. Whereas Brelaz invokes a saturation degree (i.e. a count of the differently coloured neighbours), we assess the impact of colouring a particular node on those yet to be coloured (a count of the neighbouring as yet uncoloured nodes).

Algorithm 1 GC heuristic colouring

```

while not finished do
    determine node with highest impact
    colour greedily
end while

```

This process iterates until all nodes are coloured or else an “uncolourable” node is found, in which case the algorithm fails.

1) *Spectrum Access Etiquette*: Finally, we need an etiquette for getting access to the environment, and we assume the etiquette first presented in our earlier work [9]. 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 (see 2 for an illustration); constraints (discussed in detail in Sec. III-D) are generated to protect the receivers from harmful levels of interference, and if the new caller can find a suitable channel then it can start the call.




Fig. 2. Spectrum Hole Scenario (adapted from [23])

If, however, the spectrum is too congested and a channel is not available, then we apply a local reassignment etiquette to force some devices to “shift” their assigned channels. Initially the transmitting device for the new communication broadcasts on the *common signalling control channel* (CSCC) its intent to make a call. The new receiver (and any other currently receiving devices who can hear the new transmitter) generate constraints taking the new transmission into account. We will call this the *core problem*. In the case that the core problem is *unrestrictive* (i.e. assignments for existing devices do not need to be altered to fit a new assignment), then the device sets up a call on a suitably selected block of spectrum. Otherwise, a *reassignment group* is formed with the union of the transmitters in scope of these core problem constraints. Peripheral constraints now need to be generated on the receivers of the reassignment group transmitters. These restrict the assignments transmitters in the reassignment group may take due to the interference effects of and on calls beyond their range. If an existing transmitter is to change frequency block, then it must not contribute excessive interference to other nodes (nor allow excessive interference on its own call). The new transmitter may now suggest reassignments to the transmitters in the core problem. If no such reassignment is found, then the call is then refused. If spectrum can be found then it continues until completed, after which the spectrum is released.

2) *Simulation*: We simulate a spectrum commons in which 100 devices are randomly scattered (with a uniform distribution) across a 500m by 500m area, representing a large environment. Each device is represented as an object, which maintains details of its own location (x,y coordinates) and communications. Another object simulates the CSCC, and devices broadcasts on it their transmit power, busy state,

operating channel(s), etc.. We also assume these messages can be used to determine the predicted loss between devices. Radio signal attenuation follows a simple inverse power law: $P_r = \frac{P_t}{d^\alpha}$ in which P_r is the power received at a node; P_t is the transmit power of the node concerned; d is the straight-line distance between the transmitting and receiving nodes; and α is an attenuation constant, here set to 4.

In each trial the nodes positions are randomised, and in turn each node (if idle) attempts to establish a communication on a randomly selected service with another nearby node. If the node can find an assignment of free spectrum then it will initiate a communication session, otherwise it will attempt to bargain with its active neighbours using the spectrum access etiquette described earlier. When all nodes have been considered, the trial stops. Each experiment consists of 100 trials. In each trial, we measure both the number of calls and the spectrum usage. We define *spectrum use* as the sum of channels involved in successful calls. Thus if three nodes are transmitting, one on each service, then the total spectrum utility would be 8 (1 + 2 + 5). This same experimental setup will be used in all the scenarios that follow, where we will gradually introduce cognitive devices into the population.

IV. SCENARIO 1: EXTENDING DEVICE CAPABILITY

We now consider the gradual introduction of more flexible radios into our population of legacy devices. We want to measure the impact of these new devices as their percentage increases and that of the legacy devices falls. Do the cognitive devices benefit immediately, or do they need a critical mass before the benefits appear? Does the introduction of cognitive devices have a negative impact on the legacy devices?

A. Level 1 Cognitive Nodes

We start by introducing cognitive nodes that have frequency agile front-ends spanning a wider frequency range than the legacy nodes – a 100% increase (20 channels as opposed to 10 for the legacy devices). In practice, this could be achieved by the purchase of new devices replacing old ones, or by firmware updates to existing devices. The new devices also have power control capabilities akin to the outer loop power control necessary for UMTS [24]. Each node uses the least power necessary to achieve the target SNR be achieved (we apply a safety margin of 3dB - that is external interference is limited to 3dBm above the required SNR).

B. Impact Measurement

Our first measurement will be simply the total number of calls achieved, for each population mix, for each type of call. However, since we have different services, it is possible that one type of device may appear to be disadvantaged, by enabling more calls on service C (5 channels) but achieving fewer calls in total. Thus we also measure the total spectrum use for each population mix.

Measuring the total calls does not give an easy measure of the impact on individual devices, since the ratio of cognitive to legacy device changes over time, so we need a way of

determining whether or not the different devices are equally successful. We start with a null hypothesis that all devices are equally capable of making successful call attempts. In that case, the proportion of expected successful calls would, on average, be determined by the proportion of the types of devices in the network. If 10% of the nodes are now ‘cognitive nodes, we would expect that 0.1×0.1 of the successful calls to occur between cognitive nodes, and 0.9×0.9 of the successful calls to occur between legacy nodes. Similarly we would expect 0.1×0.9 of the calls to originate and terminate between nodes of mixed capability. The full set of expected proportions of call-setups for different proportions of cognitive devices is detailed in Table II (OITl means that the device originating the call is legacy, as is the device terminating it; OITc means that the device originating the call is legacy, but that the device terminating the call is cognitive, and so on.) This then gives us a second metric, in that we can measure the deviation from this expected performance for each type of call.

Although that second metric measures the relative success of the different types of radio, it may still obscure the impact of the cognitive devices. A user of a legacy radio may not care how much call success it gets relative to other devices; instead, the user may be interested in whether or not there is an impact on the absolute performance of the device over time - that is, as other users migrate to cognitive radios, does the success rate of the legacy device drop? Thus we also measure the the success rate for each type: i.e. the number of successful calls for that type divided by the numbers of attempts for that type.

TABLE II
EXPECTED PROPORTION OF CALLS SETUP

Reconfigurability (%)	OITl	OITc	OcTl	OcTc
0	1	0	0	0
10	0.81	0.09	0.09	0.01
20	0.64	0.16	0.16	0.04
30	0.49	0.21	0.21	0.9
40	0.36	0.24	0.24	0.16
50	0.25	0.25	0.25	0.25
60	0.16	0.24	0.24	0.36
70	0.9	0.21	0.21	0.49
80	0.04	0.16	0.16	0.64
90	0.01	0.9	0.9	0.81
100	0	0	0	1

C. Experimental Results 1

We consider 10 new population mixes, consisting of 10% cognitive devices, 20%, and so on up to 100% cognitive devices. For each mix, as before, we run 100 trials and record the average numbers.

In Fig. 3 we plot the total spectrum use and the number of calls of each type. The first thing to note is that there is a steady increase in spectrum use, from an average of 26.6 channels for 100% legacy devices to 74.4 channels for 100% cognitive devices as they are more successful in establishing calls using services B and C. As we would hope, gradually introducing cognitive devices into the population increases the

spectrum efficiency. For the total number of calls, we see a rise from 14.04 to 30.9. The number of calls between legacy devices steadily declines, while the number of calls between cognitive devices increases more rapidly. The curves reflecting the number of communications achieved by mixed links (originated by a legacy device and terminated by a cognitive device - OITc - or vice versa - OcTI) are tightly coupled (with OcTI always slightly higher due to the introduction of variable powers).




Fig. 3. Communications Achieved (Legacy vs. Level 1 Cognitive)




Fig. 4. Deviation between Achieved and Expected (Legacy vs. Level 1 Cognitive)

In Fig. 4 we plot the deviation from expected behaviour for the four types of call as the number of cognitive nodes in the network increases. We can see that there is an immediate benefit for the cognitive devices, getting a significantly larger share of the number of calls than expected even when there are only 10% cognitive devices in the population. Gradually, as the population becomes saturated with cognitive devices, this benefit levels out.

In Fig. 5 we attempt to quantify the “experience” by the user of either a legacy or a cognitive device. We measure the number of devices of each type involved in calls and then for each level of network penetration divide by the proportion of




Fig. 5. Apparent Success Rate for Device Types (Legacy vs. Level 1 Cognitive)

such devices, for example, $((OITI + OITc + OITc + OcTI) / \text{\#legacy devices})$. This gives us a measure of the calls achieved by each kind of device weighted by the proportion of such devices in the network. What we see is that for a legacy device user, his/her experience is not significantly impaired as the number of cognitive devices increase; in fact due to the flexibility of the cognitive devices the legacy users experience actually seems to improve.




Fig. 6. Success Rate for Each Call Type

In Fig 6, we show the success rate for each type of call as the population changes. The most notable feature is that there is no negative impact on the legacy devices: they still continue to achieve the same or even higher success rates. The introduction of the cognitive devices does appear to be beneficial all round. The ability of the cognitive devices to make more efficient use of the spectrum means that sufficient space is still available for the legacy devices to make their calls.

V. SCENARIO 2: IMPROVING THE INTERFERENCE MODELS

In this second scenario we further add to the cognitive elements of the radio, by improving the interference model used by the devices. In particular, we introduce a constraint

model which represents the cumulative effect of multiple sources of interference, and relates it to the SINR requirement. The model was first described in [9], where it was shown to be more efficient than the standard graph models in terms of spectrum use.

A. Level 2 Cognitive Nodes

Consider the situation shown in Fig. 7, 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 co-channel with T). The task is then to assign channels to T and G such that R's SINR is greater than the defined service level. That is, $P/(\sum P_i + N) > \theta$, 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 SINR 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.

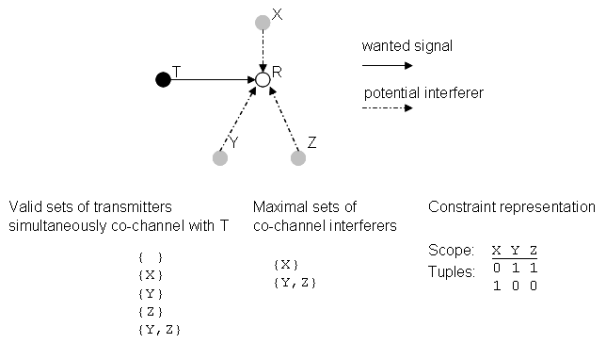


Fig. 7. Tuples in a Non Binary Constraint

In Fig. 7 the three other active transmitters whose signals could interfere with R are labelled X, Y and Z. 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. 7. 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 (very) large number of table entries.

We assume all the other features described in Scenario 1. For the spectrum access etiquette, the constraint information being exchanged between the cognitive devices is now more complicated. The legacy nodes are unable to generate or process the new constraint model. If, in the course of negotiation for spectrum, a legacy originating device requires constraints from a cognitive node, then this node will detect the capabilities of the legacy node and only provide the simpler constraint. Otherwise, exactly the same scenario in terms of device numbers, transmit powers, deployments, numbers of experimental runs etc. is used.

B. Experimental Results 2




Fig. 8. Communications Achieved (Legacy vs. Level 2 Cognitive)




Fig. 9. Deviation between Achieved and Expected (Legacy vs. Level 2 Cognitive)

Figs. 8 and 9 shows a set of results which summarises the outcome. Recall again the cognitive nodes have frequency agile front ends and adaptive power capabilities as well as the




Fig. 10. Apparent Success Rate for Device Types (Legacy vs. Level 2 Cognitive)




Fig. 11. Success Rate for Each Call Type

ability to use the more advanced receiver-centric interference-temperature constraints and adjust themselves accordingly.

In Fig. 8 the numbers of communications between cognitive nodes increases with a convex slope as the number of such devices increases in the network. The number of communications between legacy nodes (OITl) falls off more sharply than with the previous scenario (by 50% penetration of cognitive devices the legacy-legacy links achieve less than 3 on average), but there is a big jump in the number of successful communications originated by cognitive nodes to legacy nodes (peaking at 10.2) showing further reason to upgrade devices. We see that initially with only the legacy devices the numbers of calls accessing the spectrum is 13.91 (with a spectrum utility of 26.64), rising to 37.44 (90.01) when the network is fully made of the cognitive nodes.

The improvement over/impact on expected behaviour as the number of cognitive nodes in the network increases for the four categories of communication possibility as defined above, as seen in Fig. 9. The interesting issue here is that the impact on the legacy users is much greater.

Again in Fig. 10 we attempt to quantify the “experience” by the user of either a legacy or a cognitive device. What

we see is that for a legacy device user, though their initial performance is lower than before (see Fig 5), their experience is not significantly impaired.

VI. SCENARIO 3: CHANGING THE SEARCH ALGORITHM

In the third scenario we also consider the use of different search algorithm for resource allocation. In the literature on frequency channel assignment there are three main approaches: *integer programming* (and variants) [25], [26] which is not suitable for our problem since it requires a centralised solver and global knowledge of the problem; the application of heuristics derived from Graph Colouring (as discussed in previous sections); or the use of *meta-heuristic* searches [27], [28], which locally repair a candidate solution (generated greedily or randomly).

A. Local Search Algorithm

The second solution algorithm used is a relatively simple local search which attempts to *repair* an initial random solution. Assignments which cause constraints to be violated are randomly reassigned. Any such reassignment which does not worsen the overall score of constraint violations (that is improvements or sideways moves) is accepted, otherwise the potential reassignment is rejected and the previous assignment re instantiated.

Algorithm 2 Local Search

```

iterations = 0
while iterations ≤ 50 or solution not found do
  select random node
  assign random colouring
  if new colour → improvement then
    accept move
  else
    reject move
  end if
  increment iterations
end while

```

Such a local search is the basis of many heuristic combinatorial optimisation algorithms e.g. Genetic Algorithms, Simulated Annealing and Tabu Search. A fixed number of moves is made, each move being a random perturbation of the current assignment, if the new move is a solution or the number of moves is met then the search terminates (with success and failure respectively).

B. Experimental Results 3

Figs. 12 and 13 shows a set of results which summarises the outcome. Recall again the cognitive nodes have frequency agile front-ends and adaptive power capabilities as well as the ability to generate and process interference constraints and they now also use a local search algorithm when attempting to find solutions in the negotiation etiquette.

In Fig. 12 we see that initially with only the legacy devices the numbers of calls accessing the spectrum is 13.96 (with a

spectrum utility of 26.57), rising to 30.49 (76.23) when the network is fully made of the cognitive nodes. However, in this case we see a dip in the final performance as even at 90% penetration of the network with cognitive nodes, mixed node links (especially cognitive to legacy) remain effective (at 90% mix - the number of calls was 29.83 but the spectrum utility was 88.43).

While the overall numbers of calls and spectrum utility is slightly lower than with the previous scenario, in this one we see even further gains for the cognitive node originated to legacy terminated links exceeding the highest values achieved there in the 40% (10.49), 50% (11.59) and 60% (11.68) penetration levels of cognitive devices.




Fig. 12. Communications Achieved (Legacy vs. Level 3 Cognitive)




Fig. 13. Deviation between Achieved and Expected (Legacy vs. Level 3 Cognitive)

In Fig. 13 we see the gains and impacts on the numbers of communications achieved compared to that expected. We see a much more significant gain early on for the cognitive devices, though this *additional* gain in the numbers of calls achieved more than expected falls off at the numbers of cognitive devices (and hence the numbers of calls *expected*) increases (as also seen in Figs. 14). Fig. 15 shows that the *impact* on the legacy devices further reduces as the number of cognitive




Fig. 14. Apparent Success Rate for Device Types (Legacy vs. Level 3 Cognitive)




Fig. 15. Success Rate for Each Call Type

devices increases - this is due to the local search algorithm being more adept at finding solutions in a mixed constraint set.

VII. CONCLUSIONS AND FUTURE WORK

We set out to consider three questions.

- 1) Firstly, can we quantify the advantage of cognitive devices for individual users to promote early adoption? We have shown that early adopters of cognitive devices gain immediate benefits in call success rates - at each stage in each experiment, there is an incentive for users to switch from legacy devices to cognitive devices even when there are few other cognitive devices in the population.
- 2) Secondly, can we identify the general social benefit of introducing cognitive radios? We have demonstrated that as the cognitive devices are introduced, total call rates and total spectrum utilisation both increase, and continue increasing until the population is saturated.
- 3) Finally, what is the impact on existing devices as cognitive radios become more prevalent? In each of our experiments, we see that as the number of cognitive devices increase, the subjective experience of the legacy

users does not deteriorate – in fact, their call success rate increases. We believe this is because the cognitive devices are more efficient, and can adjust their own spectrum use to create space for new calls by the legacy devices. Although the share of utilised spectrum for legacy devices is dropping, this is not apparent to the user. We appear to have a win-win situation - early adopters gain an advantage, but legacy users also see their service improve.

For future work, we will investigate further aspects of device reconfigurability. There is a huge range of potential options for SDRs, and in this paper we have only considered four: the type of constraint model, the application of intelligent transmit power settings, the use of an enhanced radio front-end enabling access to a higher range of spectrum and the solution algorithm employed by the nodes when attempting to negotiate spectrum. In future work we might further consider services which allow assignments of non-contiguous blocks of spectrum (spectrum sculpting), completely different etiquettes could be implemented as their interoperability examined etc.

More generally we intend to consider the network behaviour due to a much less controlled release of software and hardware upgrades. At present we assume that certain percentages of devices in the network have exactly the same upgrade and are introduced into a network of otherwise all basic devices. We have begun work in examining what happens when the upgrades are allowed to happen more haphazardly and multiple kinds of device are allowed into the network. For example, with the current upgrade schemes described in this paper (CM, PC, RFE, SA) 16 varieties of devices are possible (some combinations are obviously more interesting than others).

There is also a lot more work to be undertaken investigating the new constraint model. We have already begun work examining their use in scenarios in which “traffic-aware” devices must co-ordinate themselves to assign channels to themselves with respect to an uneven bursty traffic model, and to explore the flexibility of these constraints in a multi-channel multi-hop routing problem. We have also begun work to examine the landscapes induced by the new constraints in the course of heuristic optimisation and found some significant differences to the structure of the landscape which a local search traverses over (particularly an increase in landscape neutrality). Better search techniques for the new constraint model are needed.

ACKNOWLEDGMENT

This work is funded by Science Foundation Ireland under Grant No. 03/CE3/1405 as part of the Centre for Telecommunications Value chain Research (CTVR).

REFERENCES

[1] J. Mitola, “Cognitive Radio: An Integrated Agent Architecture for Software Defined Radio,” Ph.D. dissertation, KTH Royal Institute of Technology, 2000.
 [2] G. Faulhaber, “Deploying Cognitive Radio: Economic, Legal and Policy Issues,” *International Journal of Communication*, vol. 2, pp. 1114–1124, 2008.

[3] L. Berlemann, “Multimode Communication Protocols Enabling Reconfigurable Radios,” *EURASIP Journal on Wireless Communications and Networking*, vol. 3, pp. 390–400, 2005.
 [4] F. Jondral, “Software Defined Radio - Basics and Evolution to Cognitive Radio,” *EURASIP Journal on Wireless Communications and Networking*, vol. 3, pp. 275–283, 2005.
 [5] R. Murty, “Software Defined Reconfigurable Radios: Smart, Agile, Cognitive and Interoperable,” *Technology@Intel Magazine*, July 2003.
 [6] I. Akyildiz, W.-L. Lee, M. Vuran, and S. Mohanty, “NeXt Generation/Dynamic Spectrum Access/Cognitive Radio Wireless Networks: A Survey,” *Computer Networks*, vol. 50, pp. 2127–2159, 2006.
 [7] L. Cao and H. Zheng, “Distributed Spectrum Allocation via Local Bargaining,” in *IEEE SECON'05*, 2005.
 [8] 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.
 [9] J. Bater, H.-P. Tan, K. Brown, and L. Doyle, “Maximising Spectrum Access to a Spectrum Commons using Interference Temperature Constraints,” in *Proc. CrownCom'07*, 2007.
 [10] FCC, “Title 47 Code of Federal Regulations: Part 15 - Radio Frequency Devices,” 1998.
 [11] 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.
 [12] B. Metzger, “Spectrum Management Technique,” in *38th National ORSA Meeting*, 1970.
 [13] M. Cheng, S. Huang, X. Huang, and W. Wu, “New Graph Model for Channel Assignment in *ad hoc* Wireless Networks,” *IEE Proceedings - Communications*, vol. 152, no. 6, pp. 1039–1045, 2005.
 [14] J. Zoellner and C. Beall, “A Breakthrough in Spectrum Conserving Frequency Assignment Technology,” *IEEE Transactions on Electromagnetic Compatibility*, vol. 19, pp. 313–319, 1977.
 [15] W. Wang and X. Liu, “List-Coloring Based Channel Allocation for Open-Spectrum Wireless Networks,” in *Proc. 62nd IEEE Vehicular Technology Conference*, 2005.
 [16] J. Riihijärvi, M. Petrova, and P. Mähönen, “Frequency Allocation for WLANs Using Graph Colouring Techniques,” in *Proceedings of the Second Annual Conference on Wireless On-demand Network Systems and Services (WONS'05)*, 2005, pp. 216–222.
 [17] J. Bater, H.-P. Tan, K. Brown, and L. Doyle, “Modelling Interference Temperature Constraints for Spectrum Access in Cognitive Radio Networks,” in *Proc. IEEE CogNet Workshop*, 2007.
 [18] W. Lehr and J. Crowcroft, “Managing Shared Access to a Spectrum Commons,” in *IEEE DySPAN'05*, 2005, pp. 134–137.
 [19] J. Kruys, “Co-existence of dissimilar wireless systems,” 2003, Cisco Systems.
 [20] S. Pollin, M. Ergen, A. Dejonghe, L. V. der Perre, F. Catthoor, I. Moerman, and A. Bahai, “Distributed Cognitive Coexistence of 802.15.4 with 802.11,” in *CrownCom'06*, 2006.
 [21] J. Eberspächer, H.-J. Vögel, and C. Bettstetter, *GSM: Switching, Protocols and Services*. Wiley, 1999.
 [22] D. Brelaz, “New Methods to Color the Vertices of a Graph,” *Communications of the ACM*, vol. 22, no. 4, pp. 251–256, 1979.
 [23] FCC, “Report and Order (FCC 03-278): Revision of Parts 2 and 15 of the Commissions Rule to Permit Unlicensed National Information Infrastructure (U-NII) Devices in the 5GHz Band,” 2003, ET Docket No. 03-122.
 [24] J. Korhonen, *Introduction to 3G Mobile Communications*. Artech House, 2001.
 [25] E. Lee and D. Lewis, *Handbook of Optimization in Telecommunications*. Springer, 2006, ch. Integer programming for telecommunications, pp. 103–128.
 [26] K. Daniels, K. Chandra, S. Liu, and S. Widhani, “Dynamic Channel Assignment with Cumulative Interference,” *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 8, no. 4, pp. 4–18, 2004.
 [27] S. Martins and C. Ribeiro, *Handbook of Optimization in Telecommunications*. Springer, 2006, ch. Metaheuristics and applications to optimization problems in telecommunications, pp. 103–128.
 [28] H. Hoos and T. Stutzle, *Stochastic Local Search: Foundations and Applications*. Morgan Kaufmann, 2005.