Auction Driven Dynamic Spectrum Allocation: Optimal Bidding, Pricing and Service Priorities for Multi-rate, Multi-Class CDMA

Virgilio Rodriguez, Klaus Moessner, Rahim Tafazolli
CCSR, The University of Surrey
Guildford, Surrey GU2 7XH, UK
v.rodriguez@surrey.ac.uk

Abstract— Dynamic spectrum allocation (DSA) seeks to exploit the variations in the loads of various radio-access networks to allocate the spectrum efficiently. Here, a spectrum manager implements DSA by periodically auctioning short-term spectrum licenses. We solve analytically the problem of the operator of a CDMA cell populated by delay-tolerant terminals operating at various data rates, on the downlink, and representing users with dissimilar "willingness to pay" (WtP). WtP is the most a user would pay for a correctly transferred information bit. The operator finds a revenue-maximising internal pricing and a service priority policy, along with a bid for spectrum. Our clear and specific analytical results apply to a wide variety of physical layer configurations. The optimal operating point can be easily obtained from the frame-success rate function. At the optimum, (with a convenient time scale) a terminal's contribution to revenues is the product of its WtP by its data rate; and the product of its WtP by its channel gain determines its service priority ("revenue per Hertz"). Assuming a second-price auction, the operator's optimal bid for a certain spectrum band equals the sum of the individual revenue contributions of the additional terminals that could be served, if the band is won.

I. INTRODUCTION

Static (fixed) spectrum allocation permanently assigns a segment of the radio frequency spectrum to a radio access networks (RAN). Static spectrum allocation can be very inefficient, in the presence of bandwidth demands that vary highly along the space dimension (from region to region) and/or along the time dimension (from hour to hour). Dynamic spectrum allocation (DSA) seeks to exploit the variations in the loads of various RAN's to allocate the spectrum efficiently. Reference [1] discusses two DSA schemes that have been recently studied. Important differences between [1] and the present work include: (i) our emphasis on a pricing-driven solution, as well as our consideration of (ii) data-transmitting terminals over CDMA, (iii) physical layer issues, (iv) the value of the service to a user ("willingness to pay").

In the present work, a "spectrum manager" implements DSA by auctioning short-term spectrum licenses. Just before the start of a DSA period, a network operator bids for spectrum, on the basis of the current state of its network. But all the awarded spectrum licenses simultaneously expire at the end of a specified short period, after which the allocation process is repeated. It is clear that a government agency could become the "spectrum manager". But there is another less obvious possibility: Spectrum owners in a given locality could create a "spectrum managing firm". They could transfer their spectrum rights to the managing firm, while maintaining ownership of this firm. And they may instruct the managing firm to use economic tools to allocate short-term spectrum rights to the original spectrum owners themselves, (and, possibly, to new communication firms that they may approve). Of course, the managing firm's profits will eventually be distributed among its owners (the original spectrum owners).

Implementing DSA via short-term spectrum licenses raises some important issues. First, what should be the "guiding principle" of the spectrum manager : revenue-maximisation, fairness, overall efficiency, etc.? Second, which market "mechanism" should be used : linear pricing, nonlinear pricing, auctions, etc.? Auctions seem a natural choice, but many formats are available, and the best choice is not obvious. Herein we assume that an auction format has been already agreed upon, such that it is optimal for each selfish participant to submit a bid that equals its true valuation of the "object". One such format is the multi-unit version of the auction proposed by [2]. In the original "second price" or Vickrey auction, "sealed bids" are submitted for an object. The highest bidder wins, but pays an amount that equals the highest *losing* bid. In practise, special provisions must be made to avoid certain types of malicious behaviour in this or any auction format. We do not discuss such measures, but refer the interested reader to a relevant discussion in [3].

We focus on the problem of the operator of a CDMA cell, populated by delay-tolerant terminals operating at various data rates, on the downlink. Given the current status of the network, the

operator must determine the maximum revenue corresponding to a given amount of spectrum, because this information determines his optimal bid. But the operator's revenues depend on its own internal pricing policies. Thus, optimal end-user prices must be found together with a bidding strategy. In [4], a parallel strand of work, we provide optimal pricing results, for a similar but simpler situation in which the spectrum manager directly sells short-term licenses at a unit price (*no auctions are used*).

There are several works that apply Vickrey-like schemes in a communication context. For instance, [5] makes a similar proposal for bandwidth allocation in a wired network (which [6] later applies to CDMA power control). But [5] assumes that each bidder can choose "bands" of an arbitrary widths, which complicates implementation and analysis. Also of interest are [7], which explores pricing issues in the downlink of a CDMA cell, and [8], which studies capacity and pricing optimisation for a communication resource subject to congestion.

II. Multi-Unit Vickrey Auction

The multi-unit second-price (Vickrey) auction when applied to spectrum allocation works as follows. The available spectrum is divided into K bands, each of width w. We assume that an operator cares about *how many* but not *which* bands he is assigned (all bands are equally good for the considered radio access technologies). Each operator submits a bid vector with K components. The first component means how much the bidder offers to pay for the first band *assigned* to him (whichever it is). The kth component means how much he offers to pay for an *additional* band if $k - 1$ bands have *already been assigned* to him. The spectrum manager receives all the bid vectors and assigns the bands as follows: first, the top overall bid (by looking at the components of all bid vectors) gets one band, the second highest bid wins the next band, and so on, until the K bands have been assigned. Notice that the overall highest and the second highest bids could be components of *the same* vector. Thus, the first several bands (possibly all) could go to the same bidder. A critical detail is that what a winner pays for a band depends on the *losing* bids of the opponents. For the first won band, a bidder will pay the highest *losing* bid submitted by the other bidders; for the next won band, he will pay the *second-highest losing* bid (excluding his own), and so on. Thus, a bidder that has won k bands, will pay the sum of the k *highest losing* bids submitted by *the other* bidders. (Ties are broken at random).

For example, let $K = 3$. A bid (b_1, b_2, b_3) means: I offer to pay b_1 if I *end up* with a *total* of one band allocated to me (I don't care which one), I offer $b_1 + b_2$ for a *total* of two bands, and I offer $b_1 + b_2 + b_3$ for all 3 bands. Suppose that only two bid vectors are submitted: $b^1 = (5, 3, 2)$ and $b^2 = (4.5, 4, 1)$. The assignment goes as follows: one of the bands goes to bidder 1 (5 is top overall bid), the next band goes to bidder 2 (second highest bid is 4.5), the last band also goes to bidder 2 (the third highest bid is 4, the second component of b^2). Since bidder 1 won only one band, he will pay the highest *losing* bid submitted by bidder 2, which was 1. Bidder 2 won two bands, and will pay the sum of the two highest *losing* bids submitted by bidder 1, that is, 3+2=5. Thus, the auctioneer will get a total of 6.

III. The motives of operators and end-users

A. Optimisation problem of a network operator

The main question the operator must answer is how much to bid for spectrum at a given DSA period. At the moment of bidding, the operator will know the number and characteristics of the terminals operating in its network, including the details of the physical communication layer (modulation, error-control coding, mode of diversity, etc). A distinguishing feature of the chosen auction format is that the bidder "best response" is to bid his "true valuation" of the object being auctioned. In our specific case, "truthful bidding" means that the first component of a bid vector should equal the maximum revenues that the operator could obtain if he gets a single band of spectrum (and nothing else). The second component should equal the *extra* revenue he would get if instead of only one band, he gets a total of two, etc. But the operator's revenues also depend on his own pricing policies. Thus, the operator must determine his own (internal) pricing policy along with the bid.

In determining an operator's end-user pricing, we neglect the competition among operators. The monopoly analysis provides some useful "bounds": it is the "best case scenario" for the operator, and the "worst case scenario" for the end-user. Additionally, this analysis may be a useful approximation of the situation often observed in practise, in which relatively few operators dominate a given region.

B. Behaviour of the terminals

We must specify the behaviour of a data terminal that can choose resources, in the presence of pricing. We focus strictly on the downlink of

a single CDMA cell. Reference [9] provides the basic physical model.

We assume a QoS index of the form $\beta_i B_i + y_i$ where (i) β_i is the monetary value to the terminal of one information bit successfully transferred (a constant for a given terminal), (ii) B_i is the (average) number of information bits the terminal has successfully transferred within a fixed length of time, say τ, and (iii) y_i is the amount of money the terminal has left after any charges and rewards are computed. This model is grounded on the micro-economic concepts of quasi-linear utility function, and partial-equilibrium analysis [10, Ch. 10].

B_i will generally depend on some physical index of quality of service (QoS), x_i (to be specified further below). When the terminal must pay $c_i(x_i)$ for QoS level x_i, it chooses x_i to maximise $\beta_i B_i(x_i) + [D_i - c_i(x_i)]$. $\beta_i B_i(x_i)$ is the "value" to the terminal of the bits it gets to transfer over the reference period (the terminal's "benefit"), and D_i is the terminal's monetary budget. D_i is just a constant for a given terminal, which limits its total expenditure. If D_i is relatively "large", it needs not be considered. Thus, the terminal chooses QoS to maximise benefits minus cost: $\beta_i B_i(x_i) - c_i(x_i)$.

IV. PHYSICAL MODEL

Before proceeding with the analysis, we specify in greater detail the physical model.

1) N is the number of terminals *receiving* data simultaneously *from* a CDMA base station (BS) (downlink operation). The BS has a total downlink power constraint of $\bar{P}$.
2) R_i bps is the data rate of terminal i
3) R_C cps is the chip rate of the channel, common to all terminals. For convenience, we set $R_C = W$, where W is the total bandwidth (spectrum) allocated to the cell.
4) $G_i = W/R_i$ is the spreading (processing) gain of terminal i.
5) Information is sent in M-bit packets carrying $L < M$ information bits.
6) $f_S(x_i)$ is the frame-success function (FSF) giving the probability of correct reception of a data packet as a function of the signal-to-interference ratio (SIR) at the receiver. Below, $f(x) := f_S(x) - f_S(0)$ replaces $f_S(x)$ to avoid certain technical problems [9]. As an example, for non-coherent FSK modulation, with packet size M=80, independent bit errors, no forward error correction, and perfect error detection, FSF is $f_S(x) = \left[1 - \frac{1}{2}\exp\left(-\frac{x}{2}\right)\right]^{80}$. However, we stress that our analysis does *not* rely on this or any specific FSF. We assume that *all we know*
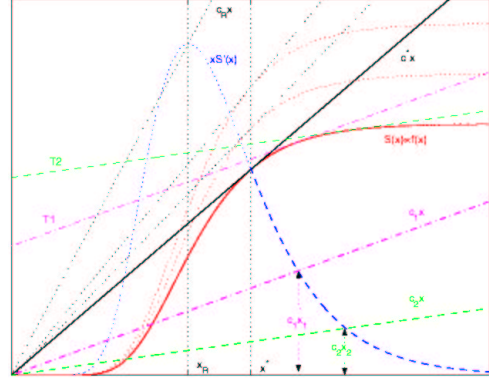

Fig. 1. From [4], pricing for revenue maximisation.

about the physical layer is that the FSF has the "S" shape shown in figure 1. The technical characterisation of an "S-curve" and some useful results are given in [11].

7) Following [7] we assume that in the downlink, the CDMA signatures retain their orthogonality, and effectively eliminate intracell interference (or that it is included as part of the random noise). Thus, the received SIR is obtained as $x_i = G_i h_i P_i / \sigma^2$ with P_i the downlink power, h_i the path gain, and σ^2 the average noise power at the receiver.
8) Packets received in error which cannot be corrected result in ideal re-transmissions until correctly received and confirmed.

A relatively simple analysis similar to that in [9] tells us that, on the average, the number of information bits successfully transferred by a terminal over the time interval τ is:

$$B_i(x_i) = \tau(L/M)R_i f(x_i) \qquad (1)$$

V. OPTIMAL SIR AND OPTIMAL PRICE

In order to determine his bid, the operator must find the maximal amount of revenues that he can obtain from the various possible amounts of spectrum he may win. Below, we focus on a single band of width w, and later generalise. It is natural for the operator to charge each terminal per unit of allocated downlink power (we consider only linear pricing). However, a price per power can be easily converted to a price per received SIR $c_i x_i$. To see this, suppose that the terminal must pay $\hat{c}$ per downlink Watt. Since the received SIR is obtained as $x_i = G_i h_i P_i / \sigma^2$, in order to enjoy an SIR of x_i, the terminal must order $(\sigma^2/G_i h_i)x_i$ Watts of downlink power, for which it must pay $(\sigma^2/G_i h_i)\hat{c}x_i$. This payment can be written as $c_i x_i$ with $c_i := (\sigma^2/G_i h_i)\hat{c}$.

The terminal chooses its received SIR x_i to maximise benefits minus cost: $\beta_i B_i(x_i) - c_i x_i$, with B_i given by eq. 1. $\beta_i B_i(x_i)$ is just a multiple of the FSF and inherits its S-shape. Thus, we must investigate how to maximise an expression of the form $S(x) - cx$, where S is some S-curve. Once the operator knows the terminal's best response to a value of c, he can choose c to maximise revenues. We perform this analysis in detail in [4], whose main results we summarise below.

First, let us think that only *one* terminal is active. Figure 1 illustrates the solution procedure. First, if the line cx lies entirely above S, except at the origin, the terminal should decline to operate, since its cost would exceed its benefit for any positive x. The largest value of c acceptable to the terminal is c^*, obtained as the slope of the unique tangent line of S that goes through the origin. For $c \leq c^*$, the maximising choice is the largest x at which the derivative of the S-curve equals c . Thus, for a given power price $\hat{c}_k$ the terminal will find the matching SIR price $c_k = (\sigma^2/G_i h_i)\hat{c}$, and will *not* operate if $c_k > c^*$. Otherwise it will choose an SIR (QoS level) x_k satisfying $S'(x_k) = c_k$ (e.g., in fig. 1 $T1$, the tangent of S at x_1, is parallel to $c_1 x$). Then, the resulting operator's revenue is $c_k x_k \equiv x_k S'(x_k)$.

As shown by figure 1, the graph $xS'(x)$ (revenues) has a single-peak at x_R, and crosses $S(x)$ at a unique positive value, x^*, precisely the point at which the tangent c^*x meets S. With the constraint $c \leq c^*$, the curve $xS'(x)$ (revenues) is maximised at $x = x^*$ corresponding to $c = c^*$ resulting in operator's revenues of $c^*x^* \equiv x^*S'(x^*) = S(x^*)$.

As long as the operator can charge an individual price to each terminal, the single-terminal analysis generalises easily. The operator will choose for terminal j a price for downlink power, such that its SIR price is c_j^*, the slope of the only tangent to S_j that goes through the origin. Since the terminals share an identical FSF, f, then each S_j is a multiple of the common f. It is easy to see and prove that replacing S_j with a multiple of S_j will change c_j^*, but will have *no* effect on the value of x_j^*. Thus, the terminals will choose an identical SIR x^*.

VI. SERVICE PRIORITY AND BIDDING

A. Power constraint

Section V tells us that as long as the terminals have a common FSF, each terminal will end up operating at the SIR x^* (but each paying an individual price c_i^* for that SIR). With bandwidth w, the necessary downlink power is obtained as:

$$\frac{w}{R_i}\frac{h_i P_i}{\sigma^2} = x^* \Rightarrow P_i^* = \frac{\sigma^2 x^*}{\bar{w}}\frac{R_i}{h_i} \qquad (2)$$

The power constraint $\sum_{i=1}^{N} P_i^* \leq \bar{P}$ can be written as:

$$\sum_{i=1}^{N} \frac{R_i}{h_i} \leq \frac{\bar{P}/\sigma^2}{x^*} w := \bar{w} \qquad (3)$$

We can think of $\bar{w}$ as the "effective bandwidth", which is the original amount amplified by the the available "rise over thermal" (ROT) per unit of required SIR, with ROT defined as $\bar{P}/\sigma^2$. Likewise, the ratio R_i/h_i can be thought of as the amount of "effective" bandwidth consumed by terminal i.

B. Service priority: revenue per Hertz

Constraint 3 may *not* be satisfied with bandwidth w. Thus, the operator needs a "service priority". Such policy is obtained by comparing the revenues provided by a terminal (if served) to the amount of resources it consumes. Terminal i, if served, will contribute revenues of $S_i(x^*) \equiv \tau(L/M)f(x^*)\beta_i R_i := \bar{\tau}\beta_i R_i$. We can think of $\bar{\tau} := \tau(L/M)f(x^*)$ as the "net" or "effective" time of operation, and $\beta_i R_i$ as the revenues from i per (effective) time unit (we can choose the time scale so that $\bar{\tau} = 1$). By dividing $\beta_i R_i$ by R_i/h_i we obtain terminal i's contribution to revenues per unit of (effective) bandwidth: $\beta_i h_i$. Now, with the terminals' labels such that $\beta_1 h_1 \geq \cdots \geq \beta_N h_N$, our service criterion is simple and clear: serve terminals 1 through I_1^*, with I_1^* the largest index such that,

$$\sum_{i=1}^{I_1^*} \frac{R_i}{h_i} \leq \bar{w} \qquad (4)$$

C. Bidding

The preceding subsection tells us immediately what the operator should offer for a single band (the first component of the bid vector), namely $\sum_{i=1}^{I_1^*} \beta_i R_i$, for a convenient time scale. To know how much to bid for an additional band, the key is to determine the additional terminals that can be served, which would tell us the additional revenue brought by the band. Assuming that the chip rate can be adjusted to match a larger bandwidth (this is *not* strictly necessary), we can multiply the right-hand side of constraint 4 by two, and obtain I_2^* as the largest index that can satisfy the new constraint, meaning that terminals $I_1^* + 1$ through I_2^* could now be served. Likewise, we can determine that terminals $I_2^* + 1$ through I_3^* could additionally be served with a third band, and so on. Then, the jth component of the bid has the simple form (with $I_0^* := 0$)

$$\sum_{i=I_{j-1}^*+1}^{I_j^*} \beta_i R_i \qquad (5)$$

and represents the contribution to revenues of the additional terminals that can be served if the jth band is won. For example, with 3 total bands and 6 active terminals, a bid vector may look like

$$\begin{bmatrix} \beta_1 R_1 + \beta_2 R_2 + \beta_3 R_3 & \beta_4 R_4 + \beta_5 R_5 & \beta_6 R_6 \end{bmatrix}$$

VII. Discussion

We have solved the problem of a CDMA operator participating in a dynamic spectrum allocation (DSA) scheme, in which a "spectrum manager" periodically auctions short-term spectrum licenses. Delay-tolerant terminals operate at dissimilar data rates in the downlink of a CDMA cell, and each user has a distinct "willingness to pay", β_i (the most he would pay for a successfully transferred information bit). We have described a realistic business model which could support our scheme.

With conveniently chosen units, our analytical results acquire crisp form. The interests of operator and end users meet at a specific operating point: the SIR value x^*, which is easily found by drawing a tangent to the graph of the frame-success rate function (FSF). At the optimum, a served terminal's contribution to revenues is $\beta_i R_i$ (with R_i its data rate), the amount of "effective spectrum" it consumes is R_i/h_i (with h_i its path gain), and its service priority is $\beta_i h_i$ ("revenue per Hertz"). The operator's optimal bid for a (an additional) band of spectrum, takes the simple form $\sum \beta_i R_i$ with the sum covering the (additional) terminals that can be served, if the band is won.

We assume that *all we know* about the physical layer is that its FSF is some S-curve. This makes our analysis relevant to a wide variety of practical configurations, and allows us to perform link layer adjustments for revenue maximisation. For instance, if modulation schemes A & B are available, each produces its own FSF. For each S-curve, we can determine the corresponding operating point $(x_m^*, f_m(x_m^*))$, by drawing the tangent mentioned above. For a given bandwidth, *other things being equal*, the revenue provided by any served terminal is directly proportional to $f_m(x_m^*)$, and the number of terminals that can be served is *decreasing* in x_m^* (see constraint 3). Thus, if $x_A^* \leq x_B^*$ and $f_A(x_A^*) \geq f_B(x_B^*)$ the modulation scheme A should always be chosen. Otherwise the operator will alternate between A and B depending on the specific values of β_i, R_i and h_i, determined by who is active and where.

Some of the additional functionality necessitated by DSA is discussed in [1]. In particular, we would like the CDMA network to adjust its chip rate to match available bandwidth. But we can also apply our scheme with an inflexible chip rate, if it matches the basic spectrum bands being auctioned. While current networks and standards do not support DSA, the needed functionality seems well within reach, given the steady advance of technology. Before any network upgrade, a cost-benefit analysis should be performed. With our results, we can perform auction-driven DSA among CDMA operators. But this will not yield the most impressive gains, because their "loads" are likely to be highly correlated. With a UMTS and a DVB-T operator participating in a DSA scheme, [1] reports gains approaching 40%. Before we introduce a wireless broadcast operator into our scheme, we must complete an analysis similar to the present one, in order to understand its bidding behaviour. Then, we could compare our gains to those that have been reported.

References

[1] P. Leaves, K. Moessner, R. Tafazolli, D. Grandblaise, D. Bourse, R. Tonjes, and M. Breveglieri, "Dynamic spectrum allocation in composite reconfigurable wireless networks," *IEEE Comm. Magazine*, vol. 42, pp. 72–81, May 2004.

[2] W. Vickery, "Counterspeculation, auctions and competitive sealed tenders," *Journal of Finance*, vol. 16, pp. 8–37, 1961.

[3] J. McMillan, "Selling spectrum rights," *J. of Econ. Perspectives*, vol. 8, pp. 145–62, Sum 1994.

[4] V. Rodriguez, K. Moessner, and R. Tafazolli, "Market-driven dynamic spectrum allocation: Optimal end-user pricing and admission control for CDMA," *Proc. of IST Mobile & Wireless Comm. Summit*, Jun 2005. (to appear).

[5] P. Maille and B. Tuffin, "Multi-bid auctions for bandwidth allocation in communication networks," *Proc. of IEEE Infocom*, vol. 1, pp. 54–65, Mar 2004.

[6] P. Maille, "Auctioning for downlink transmission power in CDMA cellular systems," *Proc. of ACM MSWiM*, Oct 2004.

[7] P. Liu, P. Zhang, S. Jordan, and M. L. Honig, "Single-cell forward link power allocation using pricing in wireless networks," *IEEE Trans. on Wireless Comm.*, vol. 3, pp. 533–43, Mar 2004.

[8] J. K. MacKie-Mason and H. R. Varian, "Pricing congestible network resources," *IEEE J. on Selected Areas in Comm.*, vol. 13, pp. 1141–9, Sep 1995.

[9] V. Rodriguez, "Robust modeling and analysis for wireless data resource management," *IEEE WCNC*, vol. 2, pp. 717–722, 2003.

[10] H. R. Varian, *Microeconomic Analysis*. New York: W.W. Norton & Co., 3rd ed., 1992.

[11] V. Rodriguez, "An analytical foundation for resource management in wireless communications," *IEEE Globecom*, vol. 2, pp. 898–902, December 2003.