Benjamin, R., Kaya, I., & Nix, AR. (1999). Smart base stations for “dumb” time-division duplex terminals. *IEEE Communications Magazine*, 37(2), 124 - 131. [No 2]. <https://doi.org/10.1109/35.747261>

Peer reviewed version

Link to published version (if available):
[10.1109/35.747261](https://doi.org/10.1109/35.747261)

[Link to publication record in Explore Bristol Research](#)
PDF-document

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Smart Base Stations for "Dumb" Time-Division Duplex Terminals

Ralph Benjamin, Ismail Kaya, and Andrew Nix, University of Bristol

ABSTRACT Users of mobile IT systems are calling for ever higher data rates, but such mobile radios are subject to multipath fading and intersymbol interference, often calling for complex equalizers at both ends of each link. Where the links are of short range, they often use time-division duplex. This article demonstrates how this option permits much of the complexity of channel matching and equalization to be transferred from post-processing at the many sparsely used mobile terminals to a few intensively used preprocessors at the base station or central hub.

Most mobile communication networks, or mobile tail-links to fixed networks, operate via a centralized base station. With the trend to high-capacity data links and multimedia applications, a mobile link is likely to suffer multipath time spreading, fading, Doppler shifts, and other complications, which demand sophisticated channel-matching, equalization, diversity combining, and other special facilities at both the base station and mobile terminals. Both the spectral demand and the propagation problems of high-data-rate applications cause the mobile-to-base links to be mainly short-range, as typified by the fast-growing areas of microcellular systems, radio local area networks (R-LANs), and wireless local loops.

In these scenarios, the two-way propagation delay is short compared to the transmission time slot, and this time slot can generally be made less than the worst-case coherence time of the propagation medium. This scenario is well matched to time-division duplex (TDD) operation, where the uplink and downlink of a duplex channel transmit alternately, on the same frequency, with their respective transmissions separated by the propagation delay. Hence, TDD is used in systems such as CT2, DECT [1], and NTT-VJ25 [2], and is a popular option worldwide, particularly since it uses the time-bandwidth product efficiently, as well as economizing in frequency sources and avoiding concurrent transmission and reception. (Usually, TDD is then combined with time-division multiple access — TDMA — since this further economizes in frequency generation, and drastically reduces the peak power and power fluctuations at the base station transmitters.)

In this article we discuss the exploitation of a further benefit of TDD. The common frequency for the up (mobile-to-base) and down (base-to-mobile) links means that all characteristics of the propagation path are identical for both, provided the path geometry does not alter significantly within the duration of one transmit-and-receive time slot. Hence, TDD permits most of the functions of matching the mobile receivers to the instantaneously prevailing propagation conditions to be handled instead by preprocessing the signals transmitted from the base station.

The resulting savings in complexity, bulk, weight, and power consumption at the mobile terminal is of course highly desirable in its own right. However, if, say, of 1000 mobile terminals within the base station's area not more than 20 are actively engaged in communication at any one time, 20 static preprocessors can

take the place of 1000 mobile post-processors, thus also permitting a very major savings in total system cost. Indeed, when TDD is associated with time-division multiplexing (TDM), a single time-shared equalizer and diversity

combiner at the base station can serve the up- and downlinks of all the mobiles within the cell. Similar benefits can be obtained in star-connected cable networks. Several authors have recognized the theoretical possibility of two-way equalization [3, 4] and two-way diversity operation [5]. This article discusses these benefits and how they can be achieved in a practical system design.

The following section deals with multipath time dispersion and shows how matching for both the up and down links can be performed at the base station. The article then explains the crucial problem of intersymbol interference (ISI), and shows how it can be largely eliminated by adaptive equalization at the base station for both the up- and downlinks. We next discuss the power fluctuations resulting from pre-equalization at

■ **Figure 1.** A typical channel: a) geometrically distinct paths; b) means and fluctuations of the resolvable composite arrivals.

the base station transmitter, and asymmetric link budgets. The article introduces the problem of deep fades due to destructive interference, discusses how diversity operation, to mitigate such fading, can be combined with path matching and ISI cancellation, and indicates how this combined function can be handled at the base station for both the up- and down-links.

We discuss the scope for *two-way* synchronization at the base station, and the effects of imperfectly aligned frequency sources and Doppler shifts. Finally, we review the status and prospects of the work, and draw a number of conclusions.

PATH DISPERSION AND MATCHED FILTERING

MULTIPATH POWER DISPERSION

In an urban environment there are normally a large number of geometrically distinct propagation paths, which we model as 40 such paths spread over $P = 5$ symbol durations. The shortest possible path is, of course, the line of sight, if available. Longer paths will normally suffer progressively more attenuation from spreading losses, reflections, scattering, diffraction, and so on, thus justifying the limit of $P = 5$ in our scenario. The time-division structure and other system parameters will be chosen so that these paths do not change significantly *within a single up- and downlink TDM time slot*. However, they may have changed quite significantly by the time the whole TDM frame is complete and the next TDD time slot comes along. A receiver matched to the symbol duration T cannot distinguish the postulated eight paths arriving in an interval T , and so gives a single output representing their vector sum. Clearly, if one of these eight paths dominates in amplitude, the relative phase of the other seven has limited effect on the magnitude (or phase) of the output vector. However, in accordance with established practice, our model assumes that the eight vectors coincide in time and have equal amplitudes, but random phase. Thus, their combined magnitude can fluctuate widely from one TDM frame to the next, and so will at times experience deep fades. The P (in our model 5) resolvable composite arrivals will of course fade independently, and are most unlikely to experience deep fades at the same time. Hence, a matched filter, combining them, can drastically reduce fading problems.

We represent the greater losses of the longer paths by appropriately reducing the magnitudes of the eight equal-amplitude random-phase vectors, from one resolvable interval T to the next. However, because of their independent fading there will still be occasions (i.e., TDM frames) when, say, the second or even the third arrival is the strongest, rather than the first. Figure 1 illustrates a typical channel and its profile in a representative TDM frame.

Thus, each TDM frame is characterized by its own distinct effective channel profile (i.e., its impulse response with a time resolution T), represented by the amplitudes and phases of P consecutive taps on a delay line. Different channel profile types have very different effects on the channel's performance, and the parameter conventionally adopted for characterizing a channel is its *normalized root mean square (RMS) delay spread*:

$$T_{RMS} = \sqrt{\frac{\sum \{(\tau_i - \tau_a)^2 A_i^2\}}{\sum \{A_i^2\}}}$$

where A_i and τ_i are the amplitude and arrival time of a specific resolvable arrival, and τ_a is the average delay.

$$\tau_a = \sum \tau_i A_i^2 / \sum A_i^2$$

(T_{RMS} could relate to a single TDM frame, but normally it is

■ Figure 2. Model of a multipath channel and its matched filter.

defined for the average over a large number of such frames to make it independent of fading.) $T_{RMS} = 0$ would denote no spread outside a single resolvable arrival (i.e., no ISI, but also no diversity due to multiple independently fading arrivals). Hence, performance is then limited almost entirely by fading, making some message frames vulnerable to noise, and losing a few more deeply fading frames entirely.

On the other hand, for five equal resolvable arrivals (i.e., virtually *maximum* ISI, but also maximum diversity due to multiple independently fading arrivals), $T_{RMS} = 1.4$. Subject to a modest minimum signal-to-noise ratio (SNR), the performance is then limited predominantly by interference from late arrivals of preceding symbols and early arrivals of subsequent ones; that is, ISI.

MATCHED FILTERING

Thus, all the characters within a given time slot arrive as similar sets of P equally spaced arrivals. Each tap is then associated with an amplitude scaling factor A and a phase-shift ϕ (Fig. 2a). Let us now feed this signal into an inverse filter, Fig. 2b, where the tap weights are the same, the phase shifts are opposite, and the delays from each tap to the final output are complementary to the delays from the initial input to the corresponding tap in the top delay line. This is a *matched filter* [6, 7]. It produces an output of amplitude $A' = \sum |A_i|^2$ and phase $\phi_0 = 0$, with a time delay equal to the length of the delay line (i.e., to the propagation path spread $P \cdot T = T_s$). This composite output has an SNR equal to the sum of the individual SNRs of all the paths: equivalent to the total energy of all signal arrivals, focused on a defined single frequency/time resolution cell, divided by N_0 , the additive white Gaussian noise (AWGN) energy per frequency/time resolution cell (commonly — but less self-evidently — defined as the noise power per unit bandwidth), the theoretical optimum [5, 6]. The wanted pulse is then preceded and followed by a symmetrical pattern of pre- and post-pulses within the time window $\pm T_s$.

■ Figure 3. Bidirectional channel matching and equalization at the base station.

In addition to optimizing the SNR, a major advantage of matched filtering is the symmetry of its outputs in the matched condition, with the matched signal itself as the largest output, at the center tap. Hence fine synchronization only requires checking for this condition. Coarse synchronization is based on correlation with a known symbol sequence, also used for channel profiling, see the section below.

BIDIRECTIONAL PATH MATCHING

When the multipath power spread is significant, it is customary to intersperse each transmission from either terminal with a predetermined sequence of *training* or *pilot* symbols so that the receivers can deduce the A_i and ϕ_i values and set up the appropriate matched filters. However, the transmission medium and its matched filter are in fact two linear filters in cascade. Their combined transfer function is independent of the order in which the signal traverses these two filters. Hence we can limit ourselves to training signals from only the mobile terminal. The base station can then analyze these to characterize the transmission medium, and hence derive the A_i and $-\phi_i$ coefficients of the corresponding matched filter [8]. The base station will then use this matched filter in the normal way to reassemble the multiple received paths. However, it can also use a further identical filter (or in some situations even the same filter) to time spread its transmitted signals, so the propagation medium will act as its reassembly matched fil-

ter for the benefit of the mobile terminal (which is thus spared the need for its own matched filter) (Fig. 3) [9]. The downlink does, however, require a simple sync pattern for time and phase reference.

INTERSYMBOL INTERFERENCE

THE ISI PROBLEM

Unless the data rate is so low that the duration of each *bit* is substantially larger than the multipath time spread (i.e. $P = 1$), each symbol emerging from the matched filter will be pre-

ceded and followed by lower-amplitude precursor and postcursor time-domain *sidelobes*, which are liable to interfere with neighboring symbols. Thus, if a channel can be represented by a five-tap delay line, its matched filter impulse response consists of a main lobe flanked by four precursor and four postcursor time-domain sidelobes:

$$\begin{aligned}
 I_4 &= A_0 A_4 (\phi_0 - \phi_4) & \} \\
 I_3 &= A_1 A_4 (\phi_1 - \phi_4) + A_0 A_3 (\phi_0 - \phi_3) & \} \text{ precursors} \\
 I_2 &= A_2 A_4 (\phi_2 - \phi_4) + A_1 A_3 (\phi_1 - \phi_3) + A_0 A_2 (\phi_0 - \phi_2) & \} \\
 I_1 &= A_3 A_4 (\phi_3 - \phi_4) + A_2 A_3 (\phi_2 - \phi_3) + A_1 A_2 (\phi_1 - \phi_2) + A_0 A_1 (\phi_0 - \phi_1) & \} \\
 I_0 &= A_4^2 + A_3^2 + A_2^2 + A_1^2 + A_0^2 & \} \text{ main lobe} \\
 I_1 &= A_4 A_3 (\phi_4 - \phi_3) + A_3 A_2 (\phi_3 - \phi_2) + A_2 A_1 (\phi_2 - \phi_1) + A_1 A_0 (\phi_1 - \phi_0) & \} \\
 I_2 &= A_4 A_2 (\phi_4 - \phi_2) + A_3 A_1 (\phi_3 - \phi_1) + A_2 A_0 (\phi_2 - \phi_0) & \} \\
 I_3 &= A_4 A_1 (\phi_4 - \phi_1) + A_3 A_0 (\phi_3 - \phi_0) & \} \text{ postcursors} \\
 I_4 &= A_4 A_0 (\phi_4 - \phi_0) & \}
 \end{aligned} \quad (1)$$

Figure 4 shows the amplitude and phase patterns of a typical multipath channel and the corresponding matched filter output. With P resolvable paths, each symbol emerging from the matched filter suffers interference from the postcursors of the $(P - 1)$ preceding symbols and from the precursors of the $(P - 1)$ following symbols. The main lobe comprises P terms, and the side lobes comprise progressively one term less, moving away from the main lobe. The average magnitude of the individual $|A_i|^2$ main lobe terms in Eq. 1 is larger than that of the $A_i \times A_j$ sidelobe terms. Moreover, the main lobe terms all combine coherently, whereas the component terms of the sidelobes have quasi-random phases. However, precursor and postcursor time-domain lobes, spaced equally on either side of the main lobe, have equal amplitudes but opposite phases $+\phi$ and $-\phi$, the solid and dotted lines, respectively, in Fig. 4b.

The matched filter optimizes the SNR of the composite signal in relation to independent random noise contaminating the distinct arrivals of the wanted signal. However, it does not protect us from ISI. In a multipath environment of significant delay spread, this interference from preceding and following symbols, rather than noise, is normally the dominant source of errors, and adaptive filters are used primarily to combat *their* effect [8]. However, these filters are normally used:

- With continuous duplex transmission so that, after initial "training," filters at both the base and mobile terminals have to adapt continuously to a changing propagation environment, or reset periodically to a new training sequence in the transmission
- With distinct up- and downlink frequencies so that the filter settings at the base and mobile terminals cannot be identical

In this article we consider how we can benefit from the special circumstances of TDD operation, which avoids these constraints.

For an insight into the ISI problem, consider a sequence of binary phase-shift keyed (BPSK) symbols, identified as $+$ or $-$. Each of the $P - 1$ pairs of symmetrically spaced preceding and following symbols will then be scaled down in magnitude in

■ Figure 4. Profiles of the channel and its matched filter output: a) channel amplitude/delay and phasor patterns; b) channel-matched output amplitude/delay and phasor patterns.

the ratio $R = |I_j|/|I_0|$, and suffer equal and opposite phase shifts, $\pm\phi_j$. Of the four possible combinations of symbols and signs, only one will produce a vector resultant in anti-phase to the wanted signal, and one will reinforce the wanted signal (Fig. 5).

The worst possible ISI would arise when:

- All A_i are equal.
- All ϕ_i are equal.
- All symbols within $\pm(P-1)$ are of opposite sign to the wanted symbol.

This worst-case scenario is so improbable that we can disregard it. The worst practical channel is probably one where only an unfavorable combination of the two interference lobe pairs with the largest $I_j \cos \phi_j$ in the matched filter output (Figs. 4 and 5) could reduce a wanted symbol to or below the noise level, or even overwhelm it. This disrupting pattern of flanking symbols would therefore occur, at worst, once every $2^4 = 16$ symbols. (Multi-amplitude multi-phase modulation systems are, however, much more susceptible to ISI and noise.) Thus, in unfavorable channel conditions a significant minority of the symbols emerging from the matched filter could be in error as a result of ISI.

ISI CANCELLATION

We can use correlation with the training signals, whose sequence is predetermined and known, to deduce the multipath characteristics of the propagation medium, and so derive the pattern of pre- and post-pulses generated by any known pulse emerging from the matched filter. If we could correctly deduce the pattern of signal symbols, we would therefore know the interference caused by each symbol and so would be able to counter its effect. Of course, this would merely make an already error-free signal sequence even more perfect. However, we have seen that, even after matched filtering, ISI is likely to invert the sign of some of our PSK symbols. If only a small proportion of signals are received in error, there must be a fairly high probability that most or all of the symbols neighboring such an erroneous signal will be correctly received, even without the benefit of any correction. The interference they would cause to the wanted signal can thus

■ Figure 5. The four possible phase combinations of pair j of symmetrical spaced interfering symbols.

be computed and cancelled, thus correcting symbols that are in error and increasing the level of confidence in symbols that are not in error and also permitting correct cancellation of these symbols' ISI with subsequent symbols. Since the matched filter is the theoretically optimum combination of signals in the face of uncorrelated random noise, ISI cancellation must entail a (normally small) penalty in SNR.

Unlike cancellation of interference from the already identified preceding signals, cancellation of interference from yet-to-come subsequent symbols entails a delay in the detection process, and the value of these subsequent symbols can, at this stage, only be deduced with limited allowance for the effect on them of their successor symbols. We have developed a deterministic combined channel-matched filter and ISI cancellation algorithm [10] which is more stable and computationally much less demanding than the "ideal" recursive least squares

(RLS) Kalman filter [1], but gives virtually identical results. Typically it reduces a precancellation bit error rate (BER) of 25 percent, from the channel-matched filter, to 0.3 percent in 17 dB SNR.

The foregoing discussion is concerned with ISI-caused errors. In deep fades noise is likely to cause error bursts, but these can be mitigated (e.g., by diversity operation), as discussed later in this article. External interference to the base station and/or mobile receiver, unlike ISI, may well be bursty in nature. To cope with this contingency, both terminals may require error-correcting codes relating to groups of symbols spread over a longer time window than the duration of such a burst.

BIDIRECTIONAL ISI CANCELLATION

Where the nature of the propagation path and modulation scheme calls for ISI cancellation, this is normally treated as adaptive equalization, and it is performed at both terminals.

However, once again, the transfer functions of the channel and its equalizing (i.e., ISI-canceling) filter are commutative (i.e., their sequence is immaterial), so we can pre-equalize the signal transmitted from the base station, thus avoiding the need for any equalizer at the mobile terminal. In this instance,

of course, we know the sequence of symbols we have just transmitted and are about to transmit, and need not derive these from a potentially erroneous received signal. On the other hand, without phase recovery in synchronization, the downlink may lack the 3 dB coherent-detection gain and, as explained later, there may in practice be other asymmetries between the up- and downlinks.

Figure 6 is an expanded, more detailed view of the base station configuration on the right side of Fig. 3, and it illustrates the equivalence of uplink receiver processing and downlink transmitterpre-processing at the base station.

Because, as explained above, cancellation of the effect of subsequent symbols is intrinsically less than perfect, particularly on reception at the base station, a

■ Figure 6. Gain normalization, channel matching, and equalization of the uplink and equivalent preconditioning of the downlink at the base station.

small number of residual errors is likely to remain, even after ISI cancellation. However, in general, these can be eliminated by including an error detection and correction (EDC) code in the transmission protocols of the up- and downlinks (at the expense of a small reduction in the effective data rate).

TRACKING THE CHANNEL PROFILE

If desired, the base station can also use the recovered symbol sequence, rather like a known training sequence, to track any progressive changes in the channel profile in order to update the filter coefficients, say, after every 20 symbols.

Tracking the uplink path also means that, at the end of the uplink transmission, we have an up-to-date initial channel profile for pre-equalizing the downlink. However, the base station will then get no information for any further updating of the channel profile during the downlink transmission. If this became a serious limitation, we would have to shorten the TDM time slots, accepting the penalty of an increased fractional "overhead" time due to associating the fixed two-way propagation delay and synchronization time with a smaller data packet.

Particularly when moving in an urban environment, there will also be occasional abrupt and drastic changes in the propagation environment, as a previously dominant path gets cut off or a totally new path becomes dominant. Equalization can then only be restored with the next training sequence, and any signal lost in the meantime has to be recovered by error-correcting codes and/or automatic repeat requests. Here, too, a shorter time slot may be desirable to reduce the number of bits needing correction or retransmission.

CONSIDERATION OF TRANSMITTER POWER

PROBLEMS OF PEAK-TO-MEAN POWER RATIO

When the signal bits are sent in time sequence, multipath propagation presents the receiver with multiple time-displaced copies of this sequence, each with its own characteristic amplitude and phase; as we have seen, substantial processing is required to extract from this the best estimate of the true sequence originally sent. Similarly, preprocessing of the signal transmitted by the base station for optimum channel-matching and ISI cancellation will generate and superimpose further sets of multiple time-displaced copies of the true signal sequence, with consequent variations in peak power. Fortunately, in practice most commercial power amplifiers, with reasonable linearity up to 6 dB above the mean power, handle the pre-equalized base station transmitter signals quite adequately [11].

A further problem arises if engineering or interference-limiting considerations set a limit on the permissible peak power. The variable-amplitude nature of our preprocessed signal may then force us to reduce the base station's mean transmit power, and hence the mobile receiver's SNR. When a peak-power-limited system is also severely noise limited, the reduction in mean power, with perfect pre-equalization, increases the error rate (up to fivefold in the worst case, according to our simulation). However, in practice the downlink from the base station is likely to operate at somewhat higher power, and hence somewhat higher SNR, than the uplink. Fortunately, in a simple PSK system the high peaks, when the signals from all filter taps add up in phase, occur only infrequently. When they do arise, their impact on the peak power can then be controlled by hard amplitude limiting prior to final filtering. By definition, this means that, at such times, ISI cancellation will be momentarily somewhat degraded. However, if this should result in a small increase in the raw error rate, appropriate EDC at the mobile receiver can

eliminate the effect from the post-processing error rate. Simulation of an unfavorable scenario, with constant input power, shows that the fractional increase in error rate due to limiting the peak-to-mean ratio to 6 or 3 dB is negligible: 1 percent and 25 percent, respectively (with no error correction). For a given peak power, such hard limiting permits an increase in mean power.

ASYMMETRIC LINK BUDGET

As in any other hub-based system, it is economical to provide the mobile terminals with relatively simple, low-cost receivers and, more particularly, with low-cost, low-power, low-weight transmitters. The base, on the other hand, can readily afford a superior receiver, say with a 3–6 dB better noise factor, to make up for the low received uplink power. For its transmitter, the base can afford a substantially higher peak power than the mobile terminal (> 6 dB), thus more than making up for the mobile's inferior receiver and making the downlink a little more robust than the uplink. To minimize interference to third parties, the base transmitter power may then be adaptively controlled, going to this full power only when the low level of the received uplink signal indicates a maximum-loss channel.

DIVERSITY OPERATION

DIVERSITY RECEPTION

In some rare TDM frames, all the arrivals happen to be at or near a fading null at the same time. Noise-induced errors can then be of similar or indeed greater significance than ISI and, in the limit, no path-matching filter can recover the signal power, so there is also no useful signal on which an ISI-canceling filter could attempt to operate. However, the various paths arrive at the receiving antenna from different directions. Hence, in the vicinity of a null, quite small differences in antenna position can result in very large differences, in both SNR and the multipath and ISI patterns.

This principle is put to good use in *space-diversity* reception [12]. Typically this uses two or perhaps three mutually displaced antennas, so it is exceedingly unlikely that both — or all three — are in a deep null at the same time. The best solution then combines multiple space-diverse signals, just as the matched filter combines multiple time-diverse signals, weighting each signal S by its mean amplitude A to produce a combined SNR equal to the sum of the constituent SNRs. Similarly, the combination of phases and amplitudes at two orthogonal polarizations are virtually certain to be different for different paths. Consequently, a null experienced at one polarization will generally be associated with significant signal strength in the other. Hence, another form of diversity reception involves two collocated antennas of orthogonal polarization.

In addition, multiple resolvable arrivals, subject to independent Rayleigh fading, provide a form of time-of-arrival diversity for which the matched filter acts as the optimum combiner. Hence, matched filtering of each individual antenna virtually precludes a deep null, and enhances the SNR of its input to the subsequent antenna-diversity process, designed to minimize the combined effect of noise and ISI.

COMBINED PATH MATCHING, DIVERSITY RECEPTION, AND ISI CANCELLATION

The optimum ratio for combining the two channel-matched outputs from a dual-diversity antenna system depends on two factors of merit: A , the relative main-lobe amplitude (i.e., the relative SNR), and Q , the relative ratio of main-lobe to aggregate sidelobe (ISI) power. Hence, each output has to be weighted by a composite factor of merit F , involving both A

■ Figure 7. Optimum dual-diversity combining using CMF for proposed two-way operation at the base station.

and Q (Fig. 7). Evidently, even an infinite SNR cannot make up for disastrous multipath characteristics, and even a perfect path cannot compensate for zero signal power. Hence, we should expect F to be a multiplicative function of the generic type $F_i = A_i^a Q_i^q$; that is, if antenna i has individual factors of merit A_i and Q_i , these are raised to empirically derived powers a and q , respectively, and multiplied to form the optimum composite factor of merit. However, channel matching minimizes variations in SNR, and we have demonstrated by simulation that equivalent results are obtained with the simple additive formula

$$F_i = \alpha A + (1 - \alpha)Q \quad (2)$$

Extensive simulation showed the (not very critical) optimum value of α to be 0.6.

TWO-WAY DIVERSITY OPERATION

In diversity operation we can exploit antenna reciprocity: the coupling between the antenna systems at the two ends of a link is the same, irrespective of which is transmitting and which receiving. Hence, we can introduce yet a further aspect of two-way base station operation: having found the optimum conjunction of matched filtering and ISI cancellation for the individual antennas, and of optimum diversity combining of these antennas for reception, we then also use the same filter and diversity settings for preprocessing of the transmitted signal for optimum eventual reception by the mobile terminal (Fig. 7). Indeed, a simple form of two-way dual-antenna diversity operation is already in use in some DECT base stations.

The improvement in BER due to this form of diversity operation is quite dramatic; Fig. 8 shows the impressive cumulative improvements in performance with successive refinements in signal processing for two distinct multipath profiles. With a high sidelobe-to-mainlobe ratio (Fig. 8b), the time-diversity effect enhances the performance at low or moderate SNRs (i.e., when noise-limited). However, at high SNRs (i.e., when ISI-limited), the high side-lobe energy results in a higher ISI error rate. Equalization and pre-equalization are in fact not fully commutative because of the intrinsic nonlinearity of the decision process in the feedback section of the equalizer. For some rare channel profiles this impairs the relative performance of the pre-equalizer so that, at high SNR, it limits at a higher ISI-caused error rate than the extremely low error rate eventually

reached by the equalizer. With diversity operation such unfavorable channel conditions are effectively eliminated. Using channel estimation based on the correlation of a 63-symbol pseudo-noise (PN) training sequence, the simulations confirm that the channel estimation errors become negligible [8].

TWO-WAY SYNCHRONIZATION

Any TDD scheme requires accurate synchronization — even more so if the TDD is associated with TDM [13]. Normally synchronization is a one-way process, controlled by the base station, where both stations have to make an adjustment according to the propagation delay experienced. However, this too can be implemented as a two-way process, with most of the burden shouldered by the base station.

The simplest situation entails a fixed cycle, corresponding to the maximum number of active links which can be accommodated, allowing both terminals of all links to be active within the same TDD/TDM frame. In this case the base station's transmission to mobile $n - 1$ is followed immediately by a call to mobile n to start its transmission. In the sort of mobile network considered, the propagation delay between the base station and the mobile terminal is large compared to the bit duration, but small compared to the duration of a message "packet." Hence, the base station then transmits to mobile n with a fixed delay, after its callup of n , equal to the time taken by the training and data transmission from mobile n plus the maximum two-way propagation time (Fig. 9). This time allocation and synchronization scheme can be refined to accommodate data, in addition to or in place of voice, and/or to make efficient use of otherwise empty or underused time slots.

FREQUENCY OFFSET

Our TDD scheme is virtually immune to imperfect alignment of the frequency sources at the two terminals of a link, since the phase reference is reset at the beginning of each up or down time slot. Assuming that:

- The time slot plus propagation delay is equivalent to 1000 bits' duration.
- Each bit comprises 20 cycles of the RF carrier.
- We can only tolerate a phase shift of 0.2 cycles over this total interval.

Normalised RMS delay spread = 0.3
(a) Interference-to-signal ratio = -15 dB

Normalised RMS delay spread = 1.1
(b) Interference-to-signal ratio = 0.4 dB

■ : equalizer (uplink; * : pre-equalizer (downlink); 1: unprocessed channel; 2: 1 plus CMF; 3: 2 plus equalizer or pre-equalizer; 4: 3 plus dual-diversity

■ **Figure 8.** Progressive improvement in the performance of up- and downlink: a) interference-to-signal ratio (ISR) = -15 dB; b) ISR = 0.4 dB.

■ **Figure 9.** Training scheme for base, "two-way synchronization" at the base station.

The frequencies need only be aligned to $0.2/(20 \times 1000) = \pm 1$ part in 10^5 , which is easily satisfied. The Doppler shift to produce this fractional frequency offset is independent of the carrier frequency. It would require the path length to change at

$$\left(\frac{1}{10^5}\right) \left(3 \cdot 10^8 \text{ m/s}\right) \left(1 \text{ km}/10^3 \text{ m}\right) \left(3600 \text{ s/hr}\right) = 10,800 \text{ km/hr}$$

However, with systems of low fractional bandwidth and large package size, Doppler shifts can become a problem. Hence, for even greater simplicity and robustness we may use differential PSK, $+\pi$ for a 1 and zero for a 0. We are then immune to any frequency offset up to, say, 10 percent of the bit rate, albeit at the cost of a 3 dB degradation in SNR.

STATUS AND PROSPECTS

We have proved two-way operation at the base station, by extensive simulation, for:

- Matched filtering
- ISI cancellation
- Antenna diversity

We have also proved channel profile tracking at the base station during the mobile station's transmission (necessarily one-way only). We have not specifically demonstrated the (noncritical and straightforward) process of two-way synchronization. We have developed a hardware implementation of matched filtering, ISI cancellation, and, if desired, channel-profile tracking based on a 16-bit 33 MIPS Analogue Devices ADSP2181 DSP. This system is semi-online; that is, the received data is held in a buffer and then processed at less than real speed, and the transmit data is preprocessed at less than real speed and then read out of a buffer for transmission. It uses a new deterministic equalizer algorithm [10], which virtually equals the theoretical limit of performance but is substantially simpler, faster, and more stable than other published algorithms. For fully real-time operation in an operational system, we would recommend the use of a more modern and powerful DSP, together with a dedicated LSI auto-correlator.

Two-way base station operation is clearly most appropriate to new microcellular networks, self-contained mobile

short-range systems, wireless LANs or local loops, or star-connected local cable networks, where propagation delays are small, but the combination of high data rate and propagation geometry is likely to cause significant multipath problems.

CONCLUSIONS

- In a time-division duplex (and preferably also time-division multiplex) network, the base station can handle channel-matching, cancellation of inter-symbol interference, synchronization, and optimum diversity combining for both the up- and downlinks, leaving the mobile terminals merely with the requirement of transmitting a combined synchronization and training signal.
- Since the base station is subject to fewer weight, bulk, and power constraints, pre-equalization at the base station is clearly more readily acceptable than the equivalent equalization at the mobile terminal.
- Since the relevant up- and downlink functions are almost identical, the cost of such two-way operation at the base station is likely to be less than twice the one-way cost.
- Most important of all: since the number of base station links to be served is very small compared to the number of mobile terminals likely to be within its cellular area, the total system financial cost of two-way base station operation is dramatically lower than the cost of sharing the operations between the base station and the mobile terminals, particularly in demanding networks requiring adaptive optimization.
- Indeed, when, as is likely, TDD is associated with TDM, a single time-shared equalizer and diversity combiner at the base station, reset by the initial training signal for each pair of duplex time slots, can serve the up- and downlinks of all the mobiles within the cell.
- The optimum apportionment of the "link budget" is also different for the uplink and downlink, respectively.
- The scheme can be made virtually immune to frequency offsets.

REFERENCES

- [1] ETSI/GSM Rec. 05.02.1990.
- [2] T. Okada *et al.*, "VJ25 System: 19-GHz High-speed Wireless LAN System," *NTT Rev.*, vol. 9, no. 1, Jan. 1997, pp. 85-92.
- [3] D. Forney and M. V. Eyuboglu, "Combined Equalization and Coding Using Precoding," *IEEE Commun. Mag.*, Dec. 1991, pp. 25-34.

- [4] W. Zhuang and W.C. Huang, "Phase Precoding for Frequency-Selective Rayleigh and Ricean Slowly Fading Channels," *IEEE Trans. Vehic. Tech.*, vol. 46, no. 1, Feb. 1997, pp. 129-42.
- [5] P. Balaban and J. Salz, "Optimum Diversity Combining and Equalization in Digital Data Transmission with Applications to Cellular Mobile Radio-Part I: Theoretical Considerations and Part II: Numerical Results," *IEEE Trans. Commun.*, vol. 40, no. 5, May 1992, pp. 885-907.
- [6] R. Benjamin, *Modulation, Resolution and Signal Processing in Radar, Sonar and Related System*, Pergamon, 1966, pp. 151-62.
- [7] K. Watanabe, "Adaptive Matched Filter and Its significance to Anti-Multipath Fading," *IEEE ICC*, 1986, pp. 1455-59.
- [8] R. E. Zimmer and R. L. Peterson, *Introduction to Digital Communication*, New York: Maxwell-McMillan, 1992.
- [9] I. Kaya, R. Benjamin, and A. Nix, "Two-way Basestation Equalization for Mobile Multi-Media Communication," *Int'l. Wksp. Mobile Multimedia Communication*, Bristol, Apr. 23-24 1995.
- [10] I. Kaya, A. Nix, and R. Benjamin, "A Fast Computation Algorithm for High Performance Equalisers Using Channel Matched Filter," UK Provisional Patent, applied on Aug. 28th 1997.
- [11] R. Wilkinson, P. Kennington, and J. Marvill, "Power Amplification Techniques for Linear-TDMA Basestations," *Proc. IEEE GLOBECOM '92*, Orlando, FL, Dec. 1992, pp. 26-30.
- [12] J. D. Parsons, *The Mobile Radio Propagation Channel*, London: Pentec Press, 1992.
- [13] J. G. Proakis, "Adaptive Equalization for TDMA Digital Mobile Radio," *IEEE Trans. Vehic. Tech.*, vol. 40, no. 2, May 1992.

BIOGRAPHIES

RALPH BENJAMIN (DrBenjamin@aol.com) received a B.Sc., Ph.D., and D.Sc. from London University. He served in (British) Naval Scientific Service as chief of research, Admiralty Surface-Weapons Establishment, then director, Admiralty Underwater-Weapons Establishment. He then became superintending director and chief scientist, GCHQ (the British NSA). On retirement, he served as head of communications techniques and networks, NATO (SHAPE Tech Center). He is now consultant and visiting professor at Imperial and University Colleges, London, and University of Bristol, originating and supervising research projects. As a Fellow of the IEE he has served on its Council and Electronics Board, and as regional chairman. He is a Fellow of City & Guilds of London Institute and the Royal Academy of Engineering.

ISMAIL KAYA obtained his M.Sc. on power electronics and control system at Karadeniz Technical University, Trabzon, Turkey (KTU) in 1991, and his Ph.D. on high-speed equalizer and pre-equalizer implementation at the Center of Communications Research, University of Bristol, United Kingdom, in May 1998. He is currently a research assistant at KTU. He is an associate member of IEE.

ANDREW R. NIX (Andy.Nix@bris.ac.uk) received his B.Eng. and Ph.D. in electrical and electronic Engineering from the University of Bristol, then joining the academic staff of that university in 1993, working on wireless LANs, including LAURA (Local Area User Radio Access), ACTS, AWACS (ATM Wireless Access Communications System), and WINHOME (Wireless Innovations in the HOME). His current researches include propagation and system modeling for high-speed wireless data applications.