FINAL REPORT 6
PART III

The Study of Space Communications Spread Spectrum Systems

REPORT 90-3

PREPARED FOR THE DEPARTMENT OF COMMUNICATIONS

UNDER DSS CONTRACT NO. 36001-8-3528/01-SS

UPLINK SYNCHRONIZATION IN AN EHF FREQUENCY-HOPPED SATELLITE SYSTEM

Department of Electrical Engineering

Queen's University at Kingston

Kingston, Ontario, Canada

FINAL REPORT 6
PART III

The Study of Space Communications Spread Spectrum Systems

REPORT 90-3

PREPARED FOR THE DEPARTMENT OF COMMUNICATIONS

UNDER DSS CONTRACT NO. 36001-8-3528/01-SS

UPLINK SYNCHRONIZATION IN AN EHF FREQUENCY-HOPPED SATELLITE SYSTEM

Final Report - Part III
Uplink Synchronization in an EHF Frequency-hopped Satellite System
D.S.S. Contract 36001-8-3528/01-SS
A study of Space Communications Spread Spectrum Systems
Phases 5 and 6

S.J. Simmons P. Lamers

Mar 31, 1990

Contents

1 Preface

This part of the report deals with uplink synchronization in the frequency-hopped multiple-access EHF satellite system. Uplink synchronization encompasses the acquisition (and tracking) of parameters required to allow the detection of frequency-hopped M-ary FSK user transmissions at the satellite which performs frequency-dehopping and detection processing. In this report, we are only concerned with *coarse* synchronization. Coarse acquisition brings the user from the point of having no (or virtually no) signal energy detected by the satellite in the user's normal channel to having some significant energy detected; it is at this point that we declare coarse acquisition. The procedures by which fine adjustments may then be made to establish near-perfect alignment are studied in [5], [6]. Since the synchronization process will require feedback to the user, the issues of downlink format and synchronization need to be addressed, but to a lesser extent. This has been done in a previous report [1]. High-level data formats and associated protocols (.e.g. for network management) are not considered.

2 Objectives and Approach

The primary aims of this study are to present and evaluate coarse synchronization procedures with respect to the average time to acquire coarse synchronization, T_{acq}, the distribution of these times-to-sync, and the distribution of residual timing errors after sync is declared. The acquisition time of course depends on the SNR (signal to noise ratio) and SJR (signal to jamming ratio) at which the system is operating, on the jammer type, on the search strategy employed, on the format of the sync probe signals and the feedback downlink return data, and on the initial parameter uncertainties. These issues are addressed in following sections. We consider primarily multitone (MT) jammers in the presence of system background noise.

Although the sync procedures are naturally attempting to distinguish between background jamming and the presence of user sync probe signals, we will not overtly incorporate any features that attempt to identify jammed hops. This may possibly lead to a longer expected time to sync, but the resulting procedure will be simple and robust. We have

combined an analytical approach with Monte-Carlo simulation to evaluate the performance of the proposed sync procedures.

Throughout it will be necessary to modify conventional frequency- hopping synchronization procedures to account for the special conditions imposed by the satellite-based receiver, namely, the large propagation delay before downlink returns can be seen, and the inability to directly access the output of the frequency dehopping receiver.

3 Introduction

The system under consideration has been defined in a previous report [1], but a brief summary is included here for completeness. The general problem of uplink synchronization is then outlined, and previous results summarized.

3.1 System under Consideration

It is assumed that this is principally a point-to-point communications system based on circuit switching. Users communicate through a processing satellite that dehops user uplink transmissions, and creates suitably formatted TDM frames for transmission on the downlink. Call setup is assumed to be handled through some central controller. This controller may be a ground station, or may in fact be on board the satellite itself. Users send call requests to the controller using the same uplink channels that they use for the subsequent call; these requests therefore enjoy the same antijam protection as normal user transmissions. Channel assignments for the users may be fixed or reconfigurable by a central controller overseeing demand assignment. To keep things simple, we assume a fixed assignment scheme with dedicated uplink/downlink slots for each user. The implications of demand-assigned systems [7] are dealt with in [1].

Uplink users are arranged in an FDMA format, with each user employing a non-coherent M-ary FSK modulation. Users transmit one of M tones every T_s seconds to send $\log_2 M$ bits per symbol. We will confine our attention to the class of users whose data symbol rate

is less than or equal to the hopping rate $R_h = 1/T_h$ such that there are $T_s/T_h = L$ hops per user data symbol, where $L > 1$. Users hop their carrier frequencies as a group (as shown in Figure 1) over a very wide band (approximately 1 GHz). The carrier phase is not preserved across hops, but the symbol boundaries are assumed to be aligned with the hop transitions. The hopping pattern which the users' frequency synthesizers follow is produced by a pseudo-random sequence generator with a period which can be in excess of several days. The hopping rate employed (.e.g. 20 kHz) is sufficient to thwart frequency-follower jammers. The hopping rate determines the effective bandwidth of the signal, and puts a lower limit on the frequency spacing of the user's FSK tones. For example, a hopping rate of $R_h = 20$ kHz means that the FSK tones of each user must be spaced at 30 kHz to 40 kHz to be accurately discriminated at the satellite. Uplink beam sharing was a possibility considered in [1]. This beam-sharing option posed no additional fundamental problems, and simply implied longer waits for access to satellite resources. It is therefore not considered further.

The downlink follows a TDM format and the composite signal may or may not be protected by spectrum spreading. If spectrum spreading is employed, the first step at the receiver is to synchronize to this spreading sequence and despread. This is a conventional synchronization problem and so downlink synchronization is assumed to be a given (it has been briefly considered in [1]). The downlink is served by a single high-gain agile antenna that provides a narrow spot beam that hops from zone to zone. The downlink TDM format is assumed to be as in Figure 2.

One frame is composed of the intervals during which the downlink antenna beam hops to each of the coverage zones. While dwelling on a zone, a conventional sync pattern is transmitted, followed by data slots for the users, as well as a common information slot in which the satellite may transmit control messages or data which facilitate initial synchronization and tracking. Users look for the sync pattern, then "read" data from their assigned time slots. As in a normal TDMA satellite system, the sync pattern creates the time reference for the user time slots that follow.

We assume that when user A is not engaged in a call or a call request, the satellite places the data detections from user A's uplink slot into user A's downlink slot. This "loopback"

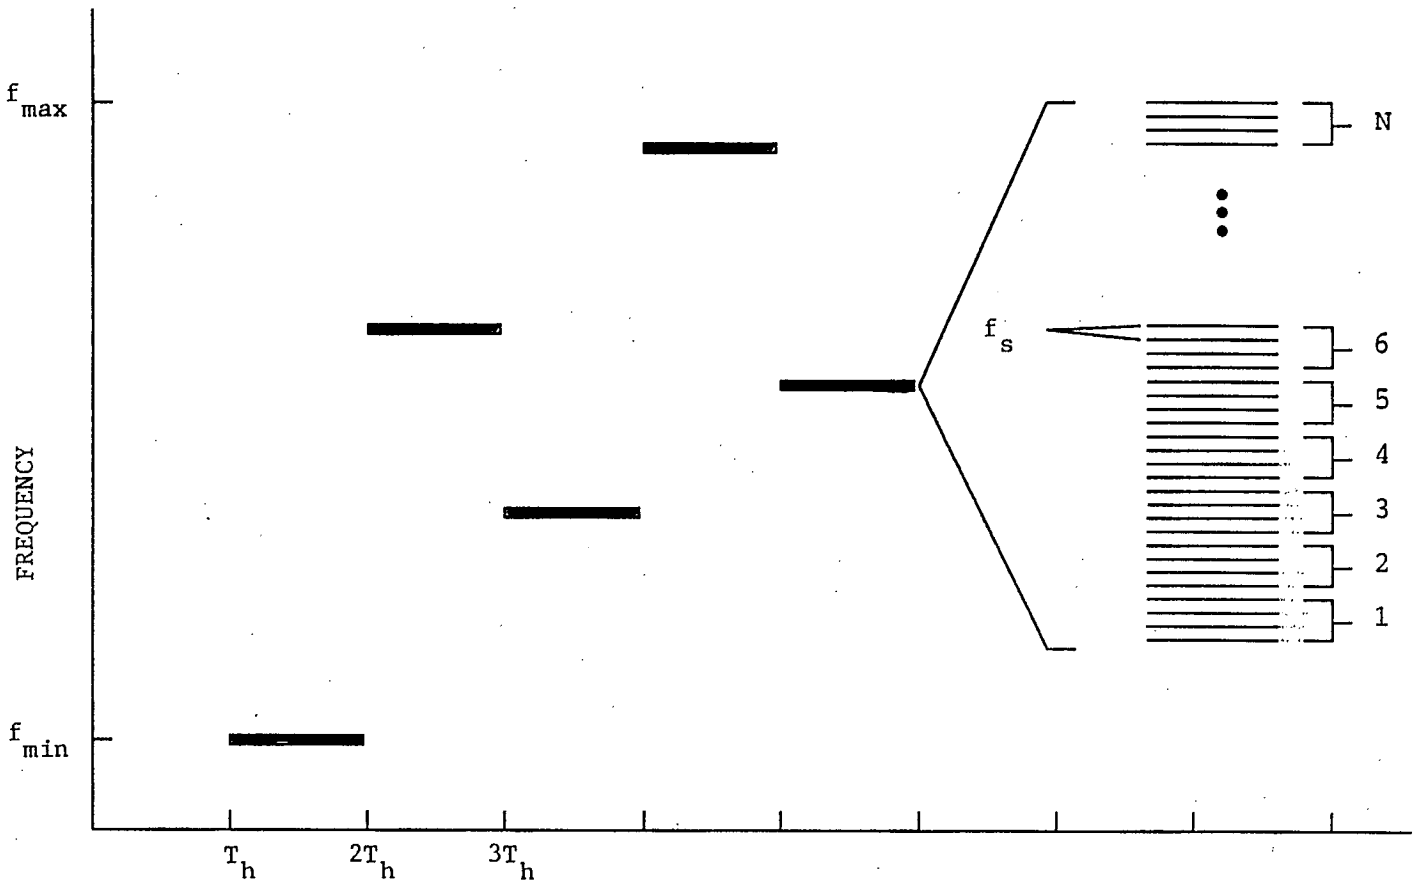

Figure 1: Uplink format for N users with 4-ary FSK

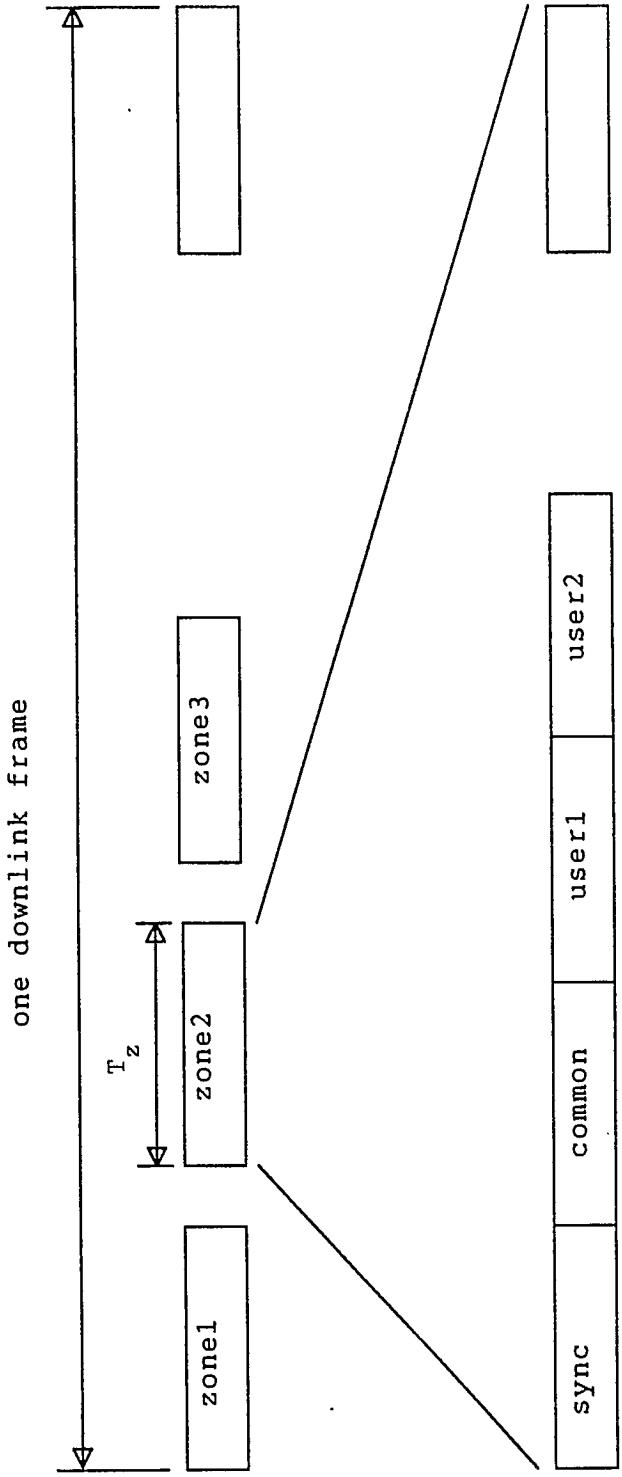

Figure 2: Downlink TDM format

mode is assumed to be the default; it provides the feedback which is crucial to initial synchronization. The type of feedback information provided for this synchronization, and the format with which it is delivered, are considered later. Once synchronization has been achieved, the network controller (which may be periodically sampling user channels to see when they have synchronized or "signed-on") can reconfigure the satellite to route subsequent transmissions to the intended call recipient. Once the call is complete, hangup codes from the initiating party will cause the central controller to terminate the call. The specific protocols by which these actions are performed do not impact on the synchronization procedure.

3.2 Uplink Synchronization Issues

Uplink synchronization refers to the process of a user adjusting his frequency hopping pattern, hop transition times, and nominal carrier center frequency in order that the satellite can detect the user transmissions. For this study, we are concerned only with coarse acquisition, that is, discovery of both the correct position in the long pseudo-random hopping sequence, and coarse carrier frequency, such that a second phase of fine adjustments may then be entered into to bring the transmissions into nominal alignment ([5],[6]). The distinction between coarse and fine synchronization here is partly a matter of degree in how much loss is encountered due to remaining misalignment, but more importantly, coarse acquisition brings the user from the point of having no (or virtually no) signal energy detected by the satellite in the user's normal channel to having some significant energy detected; it is at this point that we wish to declare coarse acquisition.

Feedback from the satellite is necessary for synchronization. We assume that when user A is not actively engaged in a call, the network controller assumes user A to be in an unsynchronized state, and configures the satellite to direct all detections from user A's uplink slot to user A's downlink slot (a default "loopback mode"). Alternatively, these "loopback" resources may be separate from the user's normal slots, and shared on a round- robin or random-access basis with other users in a demand- assigned system. The important point is that some (even if very small) fraction of satellite resources must be dedicated to every possible user in a feedback mode so that they may attempt to synchronize at a-priori unknown

(random) times.

It should be kept in mind that basic uplink synchronization must be achieved in the absence of any knowledge about the formatting of data into blocks or messages, that is, before any higher-level protocol synchronization is achieved. This implies that the synchronization process must rely on only raw symbol detection data being placed at (approximately) known positions in the downlink frames. One resulting problem is that a user first attempting to synchronize will generally not yet know his propagation delay to the satellite, and therefore cannot immediately define a one-to-one correspondence between sync probe tones sent on the uplink hops and the corresponding data returned in the downlink frames. This is discussed later.

There is one other issue of concern. We have assumed that each transmitted symbol spans several hops. The satellite may provide (diversity) processing that combines received energy over these several hops. In that case, the ground user will have to establish the correct hop at which a new symbol should be started. If there are L hops per symbol, each of L possibilities needs to be tried; it should be clear that these L trials would just represent one additional step in the searching procedure. We discuss the use of diversity in a later section.

3.3 Initial Parameter Uncertainties

Conceptually, the whole synchronization problem arises as a result of timing and frequency differences between satellite and ground equipment occurring after satellite launch. More realistically, these differences will be due to the physically distant hardware never having been synchronized by physical connection, or by drifts between equipment over the interval between the last use (when synchronization was achieved) and the current use.

The acquisitions of hopping sequence phase and coarse carrier frequency are tightly bound together. In effect, they create a two-dimensional search space over the regions of uncertainty. System implementation choices will determine the size of the uncertainty regions and ultimately the time for acquisition.

Timing uncertainty arises from differential drift between clocks on the satellite and those

in the user equipment on the ground, as well as uncertainties in the propagation delay to the satellite, t_{prop}. When a user first attempts to transmit to the satellite, it is likely that his frequency-hopper will be several hops different from that of the satellite to which he must match, and that the timing of the hop transitions will be in error with a uniform distribution over $[0, T_h]$. Note that at sync, the user's transmissions must be adjusted so that the signal arrives properly aligned at the satellite after the propagation delay t_{prop}. In order to reduce the initial uncertainty in the hopping sequence phase, the following strategy has been proposed earlier ([1]): users monitor a common slot in the downlink which contains data on the periodically-sampled internal state of the satellite's frequency hopper, and data which gives the time lag inside the satellite before that sampled state is added to the downlink data signal. If users load the sampled state into their replica sequencers, adjust for the delay from the sampling point and for downlink propagation delay, and run their sequencers ahead to adjust for uplink propagation delay, they should be in error by no more than the round-trip delay uncertainty. The only assumption needed here is that the pseudo-random sequence is produced, both in the satellite and in the ground stations, by a finite state machine (the "sequencer") realized as a clocked digital circuit. Since this is necessarily the case, there is really no assumption required. At a hopping rate of 20 kHz (50 usec per hop), a range error of 15 km translates into a hopping sequencer error of one hop. With geostationary satellites (positional drift on the order of 100 kM), and modest ground- positioning accuracy, this scheme should keep initial hopping sequence phase uncertainty to plus or minus a few hops. The range of hop phases that must be searched to acquire synchronization will be drastically reduced.

There is no need to protect this sequencer state information from discovery by an adversary. We can simply ensure that in addition to the sequencer state, users need to load a secret key that completes the true state of the sequencer. This secret key must be distributed by another secure means. The approach is described in more detail in [1].

The magnitude of the initial frequency uncertainty has not been specified. Presumably, it depends on the frequency synthesizers used. What is certain is that if the system is to operate at all, these synthesizers must be able to move the hopping carrier with an accuracy

on the order of a few kilohertz every T_h seconds. We will assume for the majority of the report that such ability implies that initial carrier accuracy will be comparable, so that we can ignore carrier frequency error entirely. However, since we have also looked at some of the issues if this is not the case (that is, the frequency synthesizers must be closed-loop adjusted to achieve accurate nominal center frequency), this work has been included in Appendix A. It assumes that the initial uncertainty in the nominal center frequency can be many times greater than the frequency separation between users.

3.4 Misalignment Effects and Search Strategies

Most generally, the search for synchronization occurs over a two-dimensional grid of cells having carrier frequencies separated by ΔF, and hopping sequence phases separated by a fraction of (or possibly a whole) hop ΔT (see Figure 3). The number of cells is determined according to the maximum initial uncertainty in both parameters. For the rest of the report, we assume zero error in the initial carrier frequency, so that the search region becomes one-dimensional (one column in Figure 3). A method for scanning through a range of hop sequence phases has been described in [1].

To synchronize, users will dwell for some time within a cell (transmitting signals as sync probes) and check the downlink return for valid detections. A user may declare firm acquisition, no acquisition, or probable acquisition. If a strategy uses declarations of probable acquisition, once such a declaration is made, we enter into a verification phase with increased observation time in the cell being evaluated. If verification succeeds, we declare firm acquistion. If the initial declaration of probable acquisition is false, we incur the time for the failed verification phase as a penalty. The probability of false alarm must be tailored to this time penalty to achieve the lowest possible overall time-to-firm-declaration of coarse acquisition. This final declaration must be made with great certainty as a very large time penalty will be incurred before any false synchronization is discovered.

The time spent in each cell and declarations made depend on the search strategy. There are a range of strategies from simple (count detections and compare to a fixed threshold, spending an equal amount of time in each cell) to complex (sequential probability ratio test

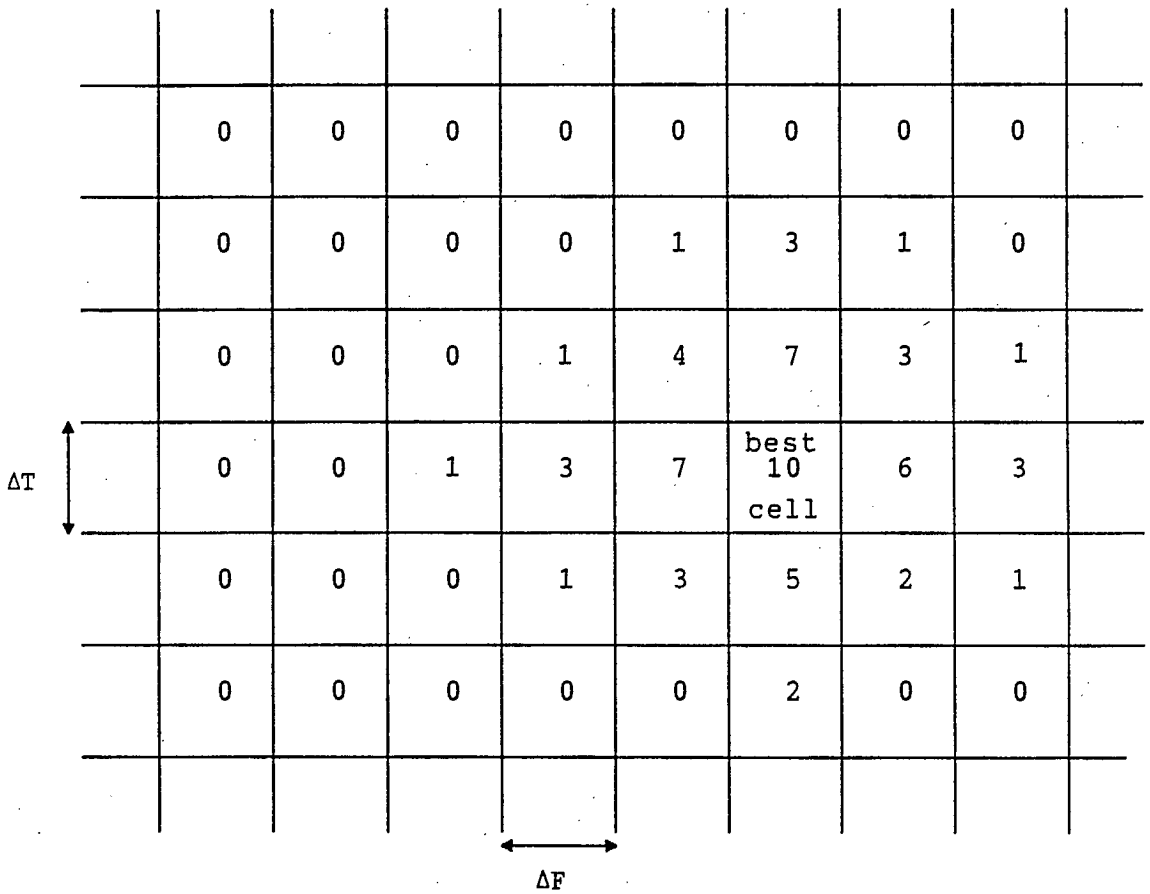

Note: Numbers indicate alignment quality in each cell.

Figure 3: 2-D search grid for coarse sync.

giving variable dwell times on each cell). The coarseness of the grid quantization (ΔT) is subject to optimization: the larger these quantities, the greater the maximum signal loss due to worst-case residual misalignment (and the longer the expected dwell time needed), but the fewer the number of cells to be searched. The choice of the cell size is strongly tied to search strategy; we comment more on this when discussing the results presented later.

It is assumed that the user transmissions are detected by a SAW based receiver that in effect performs a Fourier Transform (with additional optional windowing) on the input signal over each hop duration, producing the energy in each of the M tone bins of each user. In the presence of misalignments, the signal presented to the detector after dehopping and IF filtering is as in Figure 4 (assuming a single tone transmitted for the hop duration). After IF filtering, a portion of the sync tone will be absent with misaligned hop times (the entire tone is missing with incorrect hop sequence phases), and the remaining portion contains a frequency error. Most cells will represent complete absence of synchronization, i.e. no (or virtually no) signal energy detected by the receiver. One cell will represent best alignment or full synchronization, and the rest of the cells will represent, to varying degrees, partial synchronization. The degree of alignment is represented by the numbers assigned to the cells of Figure 3. The time errors in these cells translate directly to a loss in signal energy (assuming matched filter reception with no additional windowing) according to the following equation [2]:

$$fractional \quad loss = \frac{\sin^2 \pi \delta f (T_h - |\tau T_h|)}{(\pi \delta f T_h)^2} \tag{1}$$

where δf and τT_h are the residual frequency and time errors, respectively. Parameter τ is the timing error normalized to the hop duration T_h. For our investigation, $\delta f = 0$ so this reduces to

$$fractional \quad loss = (1 - |\tau|)^2 \tag{2}$$

With windowing in effect, the additional and the additional (multiplicative) fractional loss due to non-rectangular windowing function $w(t)$ is given by

$$\frac{\int_0^{T_c} w^2(t)\, dt}{\frac{1}{T_c}\left[\int_0^{T_c} w(t)\, dt\right]^2} \tag{3}$$

where $T_c = T_h(1 - \tau)$ is the portion of the hop that contains sync signal energy at IF.

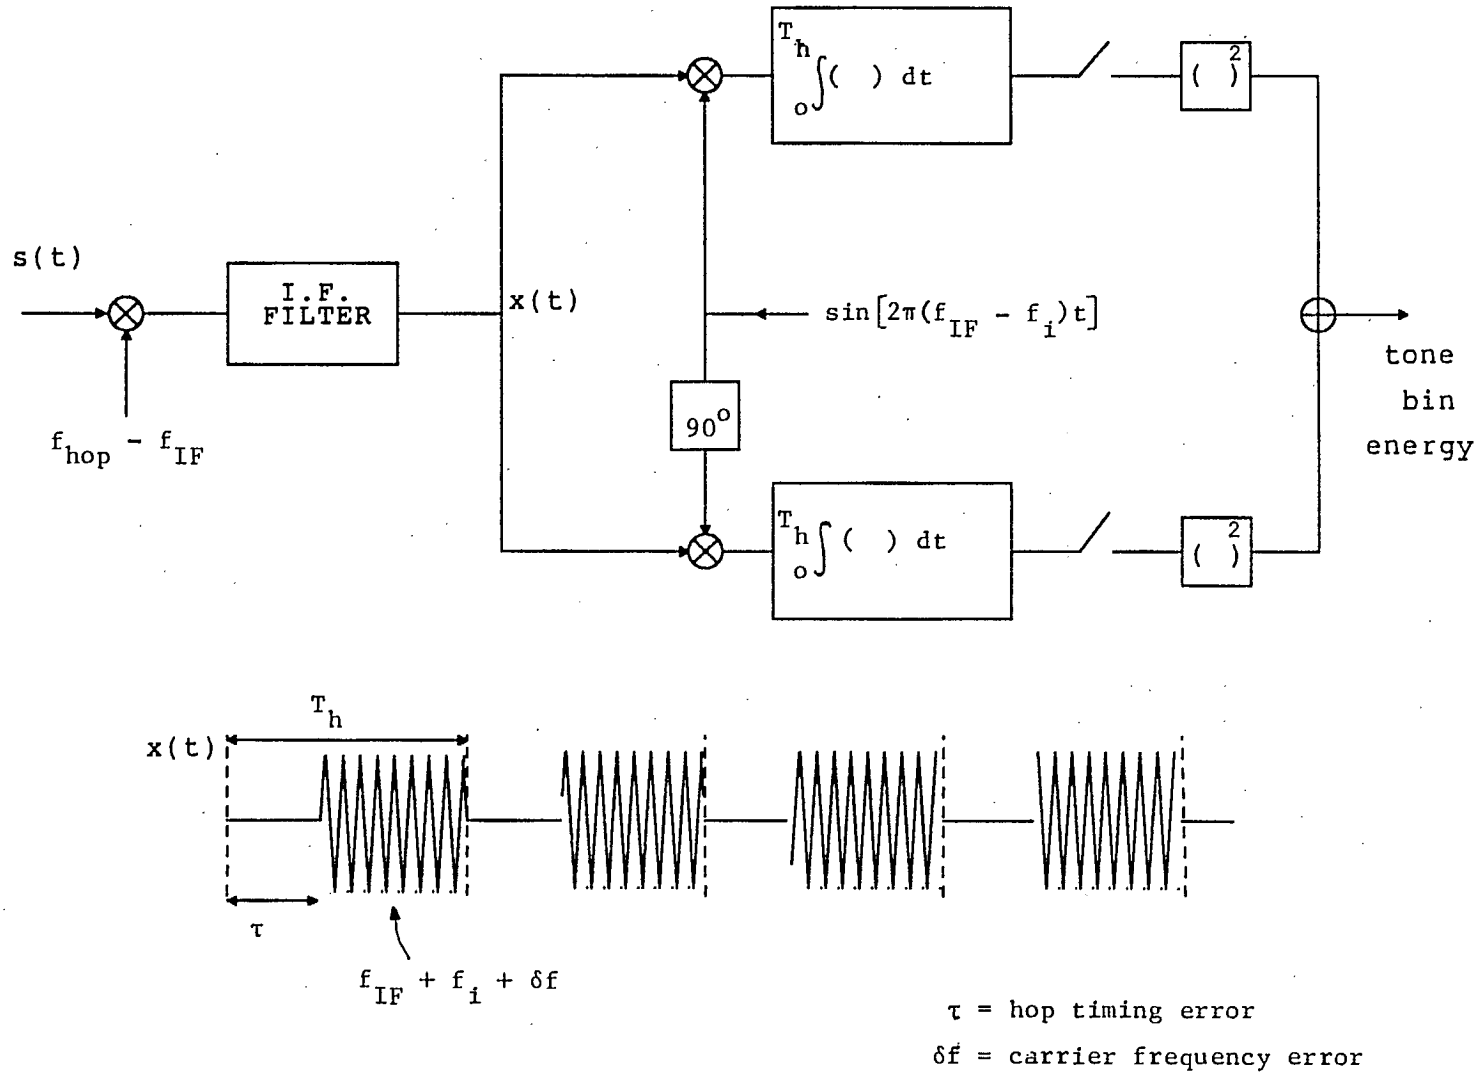

τ = hop timing error

δf = carrier frequency error

Figure 4: Detector structure and effect of misalignments.

We have evaluated the loss assuming a Kaiser-Bessel window function ([3]) with parameter $a = 1.5$ shown in Figure 5, where

$$w(t) = \frac{I_o[\pi a(\sqrt{1-(2t-1)^2})]}{I_o[\pi a]}, \quad 0 \le t \le 1 \tag{4}$$

The frequency response of this window is shown in Figure 6. In addition to signal energy loss in the sync tone bin, the timing errors also result in spectral spreading after IF, leading to spillover into adjacent tone bins. Adjacent tone bins are assumed to be separated in frequency by $1.9/T_h$. This effect is summarized in Table 1. In actual operation, we can

τ	Adjacent-bin (dB)	Sync-bin (dB)	Difference (dB)
0.00	-43.80	-5.04	38.77
0.10	-36.31	-5.22	31.10
0.20	-29.31	-5.73	23.58
0.30	-24.10	-6.75	17.34
0.40	-20.90	-8.46	12.44
0.50	-19.66	-11.05	8.62
0.60	-20.39	-14.75	5.64
0.70	-23.28	-19.93	3.34
0.80	-28.90	-27.27	1.62
0.90	-39.10	-38.64	0.46

Table 1: Windowed signal levels with timing errors

expect to sometimes declare sync in a cell even though a cell with better alignment exists. As part of the procedure, once a possible detection is declared, we will naturally step to adjacent cells to check for better alignment. Furthermore, we could consider the worst case alignment at this best cell. For example, if adjacent cells are separated by one full hop ($\Delta T = T_h$), we would assume that the best cell exists with alignment such that half the tone is missing ($\tau = 0.50$).

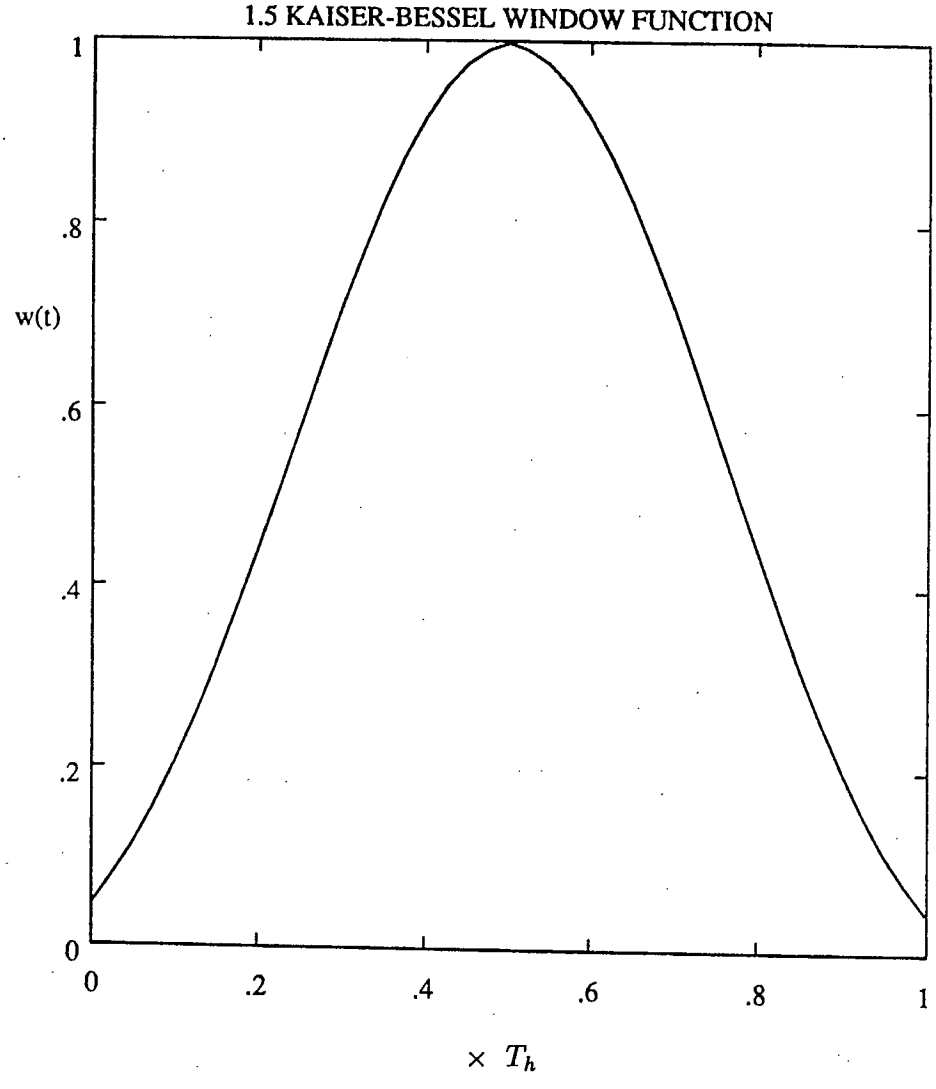

Figure 5: Receiver window function $w(t)$ normalized to T_h.

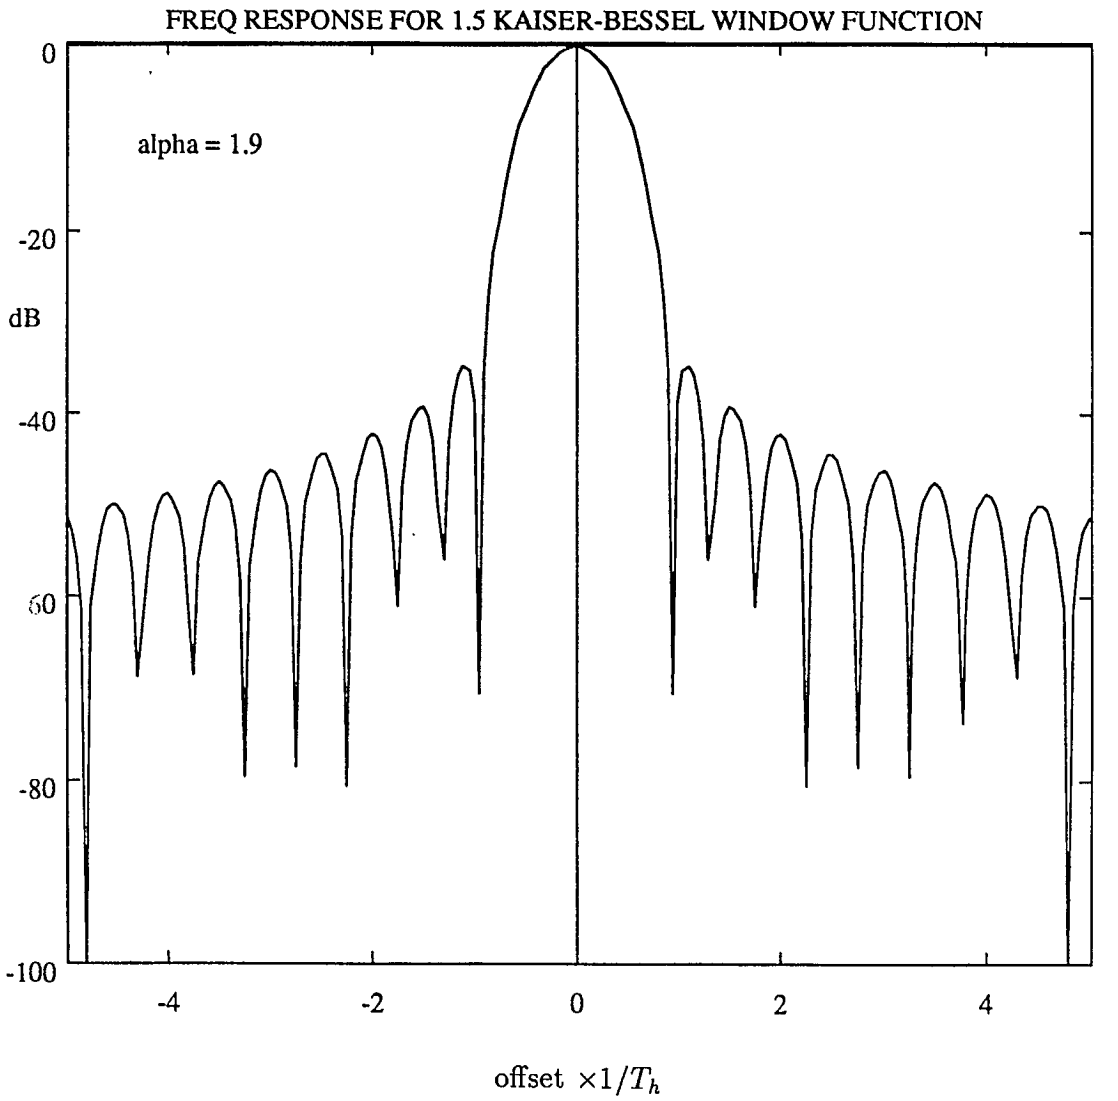

Figure 6: Frequency response for Kaiser-Bessel window.

Propagation delay places unique constraints on the search procedure used. We are unable to see the feedback from a given test signal until the full round trip propagation delay t_{prop}, which for a geostationary satellite, is in excess of 0.2 seconds. There is also buffering delay on the satellite, but this will be negligible in comparison. Two tenths of a second at a hop rate of 20 KHz amounts to 4000 hops. This does not mean that we must dwell at every cell for 4000 hops minimum. We can avoid this by "pipelining" the probe signals so that probes for evaluating many cells are issued without waiting for the return, as shown in Figure 7. By keeping a running history of the issued sync tones for each cell that is probed in turn, we can still match up the downlink returns that arrive over 0.2 seconds later and achieve a high search rate.

Our search strategy will determine how many hops we wish to dwell in each cell. If there are only C cells to be evaluated, we will have to spend at least 4000/C hops in each whether we need to or not. There is an incentive therefore to have a minimum C value that makes the search more efficient. Note that there are implications for multiple-dwell or variable-dwell searches. In these strategies, the amount of time we dwell in a cell depends on the observations made to date in that cell. The feedback delay clearly has a significant impact here. It will be necessary to break the total observation time into separate phases where these phases are separated in time by $2t_{prop}$. In each phase we send some multiple of the initial dwell time in hops, and there must be many of the same evaluations for other search cells in the pipeline for efficiency. In effect, we must be processing many search cells in parallel to avoid inefficiency due to the large t_{prop}. At the start of the search, we may be looking at a very large number of cells with a short dwell time of only tens of hops; towards the end of the search procedure we will be narrowing the search and looking at only a few cells with much longer dwell times. Clearly, the search strategy must be designed with propagation delay in mind.

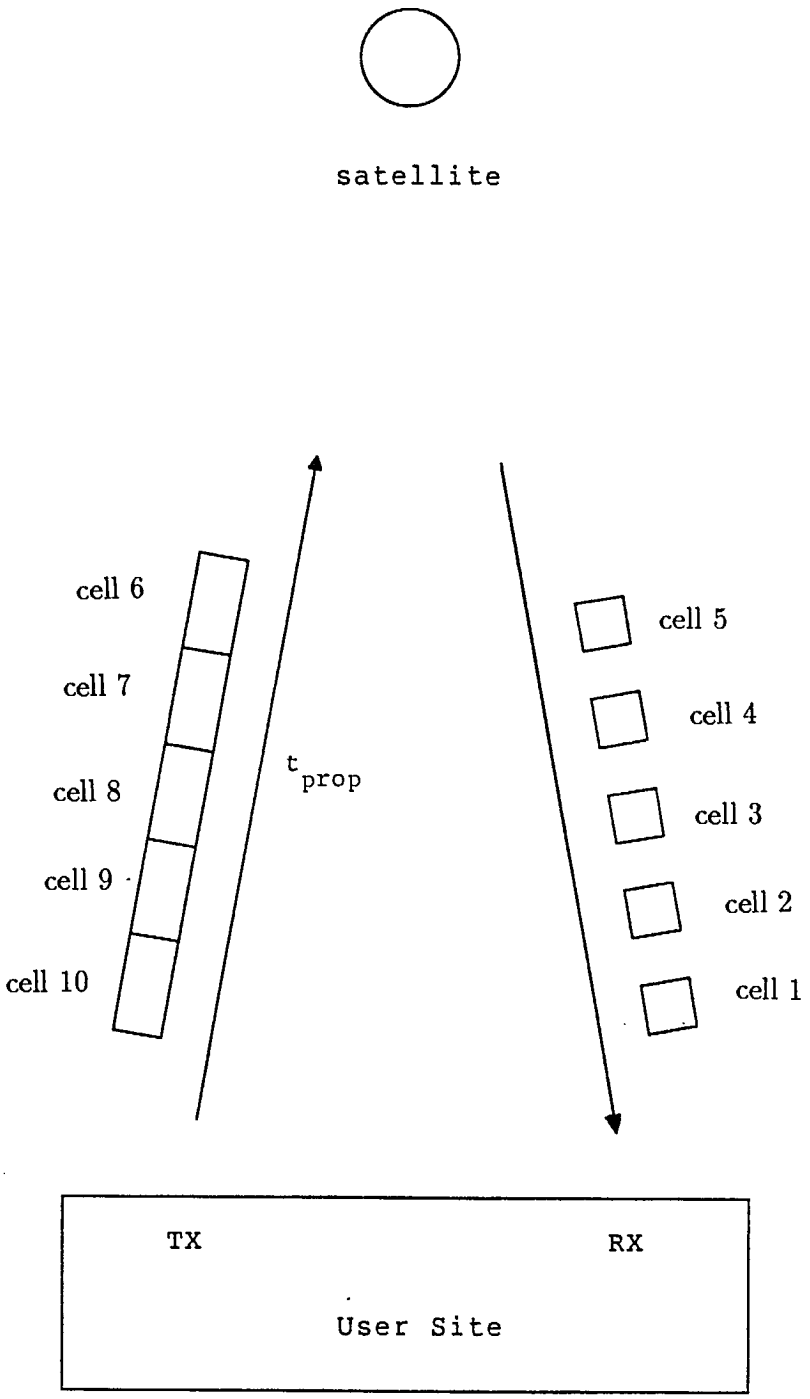

Figure 7: Pipelined sync probe signalling.

4 System Choices

In this section, we review some of the options and choices available when constructing a synchronization strategy. The discussion narrows the field and justifies the simple best "1-of-M^*" sync detection approach that is the focus of the remainder of the report.

4.1 Loopback Downlink Return Data

The downlink data returned in loopback mode depends on the type of sync probe signals we use, and the strategy we wish to detect these signals with. The number of bits we want to extract each hop depends on our detection strategy; for example, do we need access to the energy detections in many bins, or do we simply need a hard decision on one of M tone bins? The number of bits we send on the downlink will also depend on how much processing we can apportion to the satellite, and how much must be done on the ground by the user attempting synchronization. Minimization of statellite processing hardware is important, but so too is the downlink data bandwidth required for sync purposes. In fact, in a demand-assigned system, the available downlink resources for sync might be only a fraction of what is available to users actively engaged in a call. This does not necessarily mean we are overly restricted in the types of information that we wish to extract from the receiver. For example, we can take multiple-bit samples of energy detector output from the SAW for each of the M user tone bins per hop, but we may need to average these over multiple hops, or only take this for every N'th hop to match to the available downlink data rate. This latter approach is clearly undesirable as it wastes available user signal power. Whether or not diversity combining is performed on board the satellite may also impact on the format for the downlink loopback data. It also seems important to aim for a good deal of transparency, that is, to avoid requiring separate hardware and different configurations for sync purposes as opposed to normal user-to-user transmissions.

4.2 Uplink Sync Probe Signals and Detection Strategy

A critical part of the specification of the sync procedure is the type of uplink signal used as a sync probe, and the strategy with which we intend to detect the presence of this signal. The obvious approach is to send one of the M tones as the sync probe signal. This keeps the signal format identical to that needed for normal data transmissions. It is not at all obvious, however, that this is the best choice. What we are trying to do at synchronization is very different from what we are a trying to do during normal data transmission. During normal data transmission, we are attempting to send multiple bits every few hops. For synchronization, we are really only trying to discover one bit of information - whether the user signal energy is aligned or not in time and frequency - and we can take many hops to do this. In order to make this one-bit decision as quickly as possible, we would like to use a sync signal which is clearly differentiable from a background of noise and jamming, .i.e. it should look as different as possible from what we expect to see from noise and jamming (see Figure 8).

Ideally, we would like to concentrate the user signal energy in time and frequency, but we must still live within the constraints imposed by the M-ary FSK signal format, the need for frequency hopping, the limit on peak transmitter power, and the use of a SAW-based receiver. The single sync tone per hop seems to fit this bill. Spreading our available signal energy in frequency between multiple tones per hop does not appear to mesh with the objective of looking as different from jamming sources as possible. With the single tone sync probe, one precaution is necessary: the tone bin should be chosen randomly from hop to hop to thwart multi-tone jammers.

Some sync procedures may require processing to be performed on board the satellite to avoid having to send too many bits on the downlink as loopback data. To do this, the satellite may need to know which is the intended sync tone bin. This can be accomplished by selecting the sync tone bin according to the $\log_2 M$ least significant bits of the K bits that select each of the 2^K hop frequencies. In this way, at correct sync alignment, both the satellite and user agree on the intended sync tone.

There are several alternatives for detection strategies. We might consider comparing

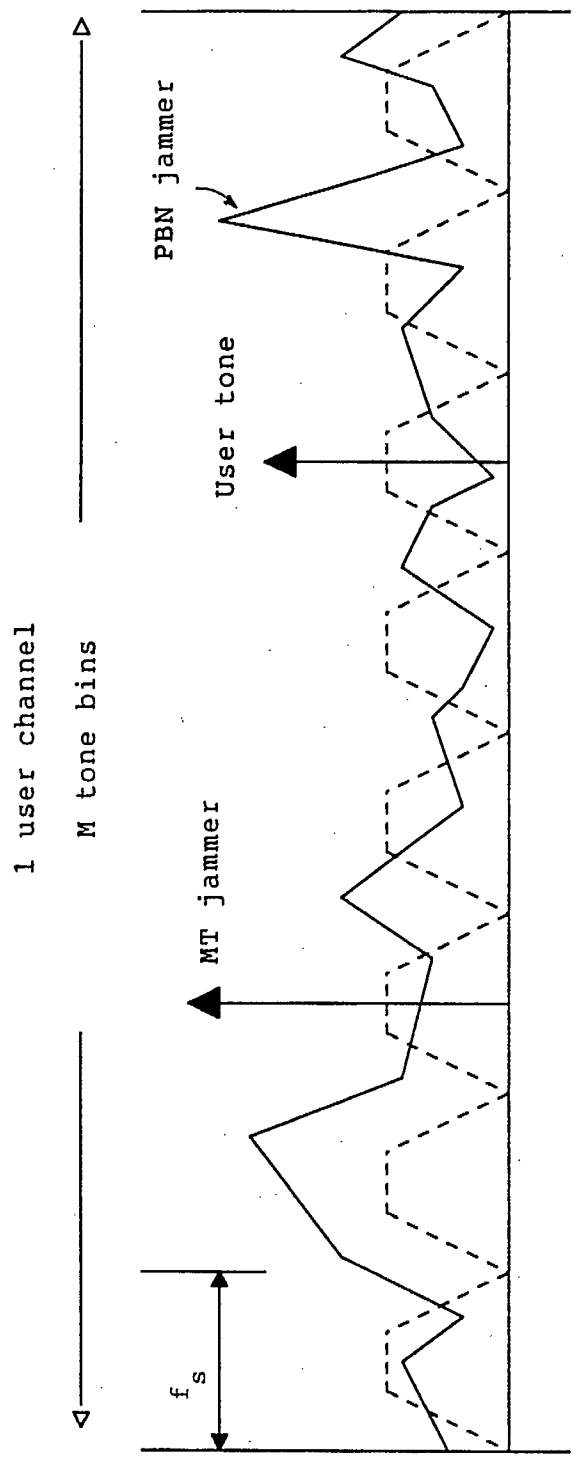

Figure 8: User sync probe signal with jamming background.

detected energy in the intended tone bin to some threshold. This energy may optionally be averaged over some number of hops with some sort of hard limit on the energy detected in the sync tone bin each hop (to prevent the presence of heavily jammed hops from skewing the result and causing a high number of false detections). The threshold must be set to give reliable detection of the presence of the sync tone even with partial alignment, while reliably rejecting background noise and wide-band jamming. A very important consideration in such a threshold scheme is the uncertainty in received signal power level at the satellite. This uncertainty will arise from unknown signal fading levels and antenna pointing errors. If we don't know what to expect for received energy in the intended tone bin, it becomes very difficult to select a useful threshold. One possibility is to add the antenna pointing angle as an additional search dimension, and set the threshold assuming some nominal fade level. This has its problems. If we are faded less than nominal, we won't be able to take full advantage of the extra signal level; if we are faded more than nominal, we may be making almost no use of the signal energy that is being received. And we still end up expanding and lengthening the search because of the additional antenna pointing-error dimension. It would be much better to avoid reliance on a threshold that does not allow for user signal power variation.

One way to accomplish this is to measure the average background jamming and noise per bin, and simply set a detection threshold equal to that level. When the user gets close to sync, the presence of user signal energy will put a positive bias on the average energy in the intended bin, and cause the detection threshold to be exceeded for a good percentage of the hops. There are two variations here. In the first, we could average the background energy detected in bins known to be unoccupied by other users. This can be done using N bins from one hop, 1 bin from each of N hops, or something in between. The energy in the sync tone bin is then averaged over an equal number of hops, N, and if larger, tone detection is declared. Note that this is not the same as declaration of sync acquisition, which may require many such tone detections.

The other variation is to instead make the detection decision on a hop by hop basis, and then count up these detections over N hops. For each hop, we can simply compare detected

energy in the intended sync bin to one other "unoccupied" bin, in effect picking the better of the two. It is not clear which of these two variations would be superior, but the latter provides an automatic hard limiting effect to prevent severely jammed hops from skewing the results, and appears to be more simple and robust than the energy-averaging scheme. We will base our sync strategies on this latter approach.

The detection scheme just outlined will produce some probability of correctly detecting the presence of the sync tone P_d (which depends on the received user signal energy plus alignment errors) and a probability P_{fd} of falsely detecting the presence of the sync tone when it is in fact not present (no user sync). A threshold will be chosen for the number of sync tone detections that need to be seen over some number of N hops before declaration of sync tone presence (i.e. acquistion or possible acquisiton) can be reliably made. This threshold will be set to produce overall probabilities of detection and false alarm that are tailored to the searching strategy to be employed. The probability of false detection per hop P_{fd} for this scheme is the probability that the energy in the intended sync tone bin exceeds the energy in one particular bin selected from the other $M-1$ tone bins when user energy is not present, and is clearly equal to 0.5. It may take many hops N to make a reliable declaration with a low enough false alarm probability. However, we can lower P_{fd} very simply. Instead of comparing the energy in the intended sync bin to that in just one of the other bins, we could compare it to all $M-1$ other bins, in effect finding the bin of highest energy. This reduces P_{fd} to $1/M$ while slightly reducing P_d. It is probable that by a judicious lowering of the detection count threshold, we can maintain a given probability of correctly declaring presence of the sync tone over the N hops observation while lowering the probability of false alarm, achieving an overall net benefit and lowering expected time-to-acquistion $\overline{T_{acq}}$.

We can get any desired P_{fd} by comparing to any number, M^*, of other bins we want. For low M, these extra "unoccupied" bins can be readily provided as common sync resource located at the edge of the user group band. There is a small practical complication with choosing other than a best 1-of-M detection approach: to avoid sending on the downlink the energies for all M bins each hop, the satellite hardware would need to make the decision on board, and only relay the bits of that decision. This requires that the satellite know the

intended sync tone bin each hop. This can be guaranteed as outlined earlier.

The number of other bins with which we compare, M^*, is subject to optimization; there will be some point at which the decrease in P_{fd} no longer provides a net benefit because the drop in P_d is too large. This optimum has been the subject of some study; specific comments are included in the discussion of the results.

5 Discussion of the Best 1-of-M^* Approach

5.1 Advantages

From the standpoint of maintaining transparency with hardware needed for normal data transmissions, the best 1-of-M^* approach suggested in the previous section is attractive. There is no need for extra satellite hardware to perform averaging or combining operations. In addition, for $M^* = M$, it generates only $\log_2 M$ bits per hop, which depending on the system implementation, may be matched to the downlink return data rate during normal data transmissions. If diversity combining over L hops is done on board the satellite, $\log_2 M$ bits per hop would represent a downlink data rate expansion by a factor of L. We might, however, simply modify the best 1-of-M^* sync procedure to use these symbol decisions rather than the hop decisions, so there would be again no downlink data rate expansion. This may, however, lead to a degradation in performance. This has not been investigated. The ramifications of on-board diversity combining are discussed below.

There are two other important benefits to a best 1-of-M^* approach. The first is that *the false alarm probability is fixed independently of jammer strategies, jamming levels, and noise levels.* The resulting sync procedure is absolutely "bullet-proof" as far as false alarms go. Jammers cannot increase the probability of false tone detection given no sync. Secondly, analysis is greatly simplified. To evaluate T_{acq}, we use the following procedure. The search procedure parameter (a threshold) is first selected. We translate worst case hop time alignment errors in the best search cell ($\Delta T/2$) into equivalent energy losses after windowing (Table 1) and apply the results of [4] to find P_d at any SNR, SJR, and jamming strategy.

Since P_{fd} is fixed independent of jammer strategy, we can then find T_{acq} by simulation given our candidate search strategy.

5.2 Diversity Considerations

Under low SJR conditions, diversity combining over L hops may be needed to achieve an acceptable error rate during normal data transmissions. A simple majority vote after best 1-of-M detection has been suggested in [4]. For completely synchronized users, a "majority-vote" diversity of $L = 9$ with $M = 8$ can achieve a symbol error rate of around 10^{-3} starting from a raw symbol error rate $P_s = 1 - P_d$ as high as 0.4 (for an unsynchronized user, the alignment errors will cause a loss of signal energy, and a corresponding detection probability P_d that is even lower). The issue arises as to whether the diversity L is fixed as part of the system specifications, or may be allowed to vary in response to changing user and/or jammer signal levels (clearly the latter is more efficient). A change in L means a change in data rate, and must be handled by appropriate protocols.

The complexity of varying L depends on whether the diversity combining is done on board the satellite or is handled by the user on the ground. Highest flexibility is achieved by leaving diversity combining to the ground user, but the cost is an expansion of downlink data rate by factor L. There are, however, problems for the synchronizing user if the combining is done on the satellite. First, majority-vote combining, while good for estimating which of M symbols was sent during normal data transmissions, is information lossy for sync purposes. Because the satellite has no way of determining which value of L a user may need, L would have to be fixed ahead of time for sync purposes. If L were fixed at the maximum value, then with good signal levels, sync might be delayed by almost a factor of L. But with poor signal levels, the information loss may be so excessive as to make sync acquisition almost impossible. Secondly, an unsynchronized user would have to search over all L possible starting positions for the diversity combining in order to discover the particular L-grouping that the satellite was using. This can lengthen acquisition by up to a further factor of L. This suggests that no diversity ($L = 1$) should be employed for sync purposes. The downlink data rate expansion in this case may be unacceptable (especially in a demand-assigned system) necessitating the

use of only every L'th hop. This again lengthens the sync acquistion time, but not as much as expected if diversity were to be used.

5.3 The Initial Correspondence Problem

When first synchronizing, a user will be uncertain as to the exact propagation delay t_{prop} to the satellite. As mentioned earlier, a range error of 15 km translates into 1 hop time uncertainty at $R_h = 20$ KHz. With a geostationary satellite having a positional uncertainty of 100 km, even if the user knows his ground position very accurately, there would be a resulting delay uncertainty of plus or minus 3 hops. What this means is that the user does not know the exact correspondence between the signals he is sending on the uplink hops and the data the satellite is returning on the downlink. The problem is depicted in Figure 9. The user transmits in one uplink frame, which the satellite detects and formats into a (delayed) downlink frame. Here we assume for clarity that the satellite is sending the user at least one bit per hop on the downlink in loopback mode. The unknown propagation delay error is represented by the horizontal offset of the frames. The user is forced to match up his k transmissions that occurred approximately $2t_{prop}$ seconds ago with each of Q possible correspondence positions in the downlink data.

Consider this example. We are using a best 1-of-M^* approach with a dwell time in each cell of 10 hops, and a propagation delay uncertainty of $Q = 6$ hops. To check the results of probing a given cell, we attempt to match the 10 intended sync tone bins (that were transmitted when probing that cell and have been stored in memory) to groups of ten contiguous detections in the downlink data, trying six different start positions for the group of ten, three of these positions delayed from nominal, and the other three advanced from nominal. Each of these $Q = 6$ positions can, however, be tested in parallel, so this does not necessarily lengthen the acquisition procedure. This does have other implications however.

First, we will be forced to dwell at least several hops in each search cell so that there is a significant chunk of consistent data with which to match to. For example, if we only send two hops per search cell, it will be quite impossible to identify when the correct cell was probed even if the sync tone was correctly detected in both hops; there will just be too many

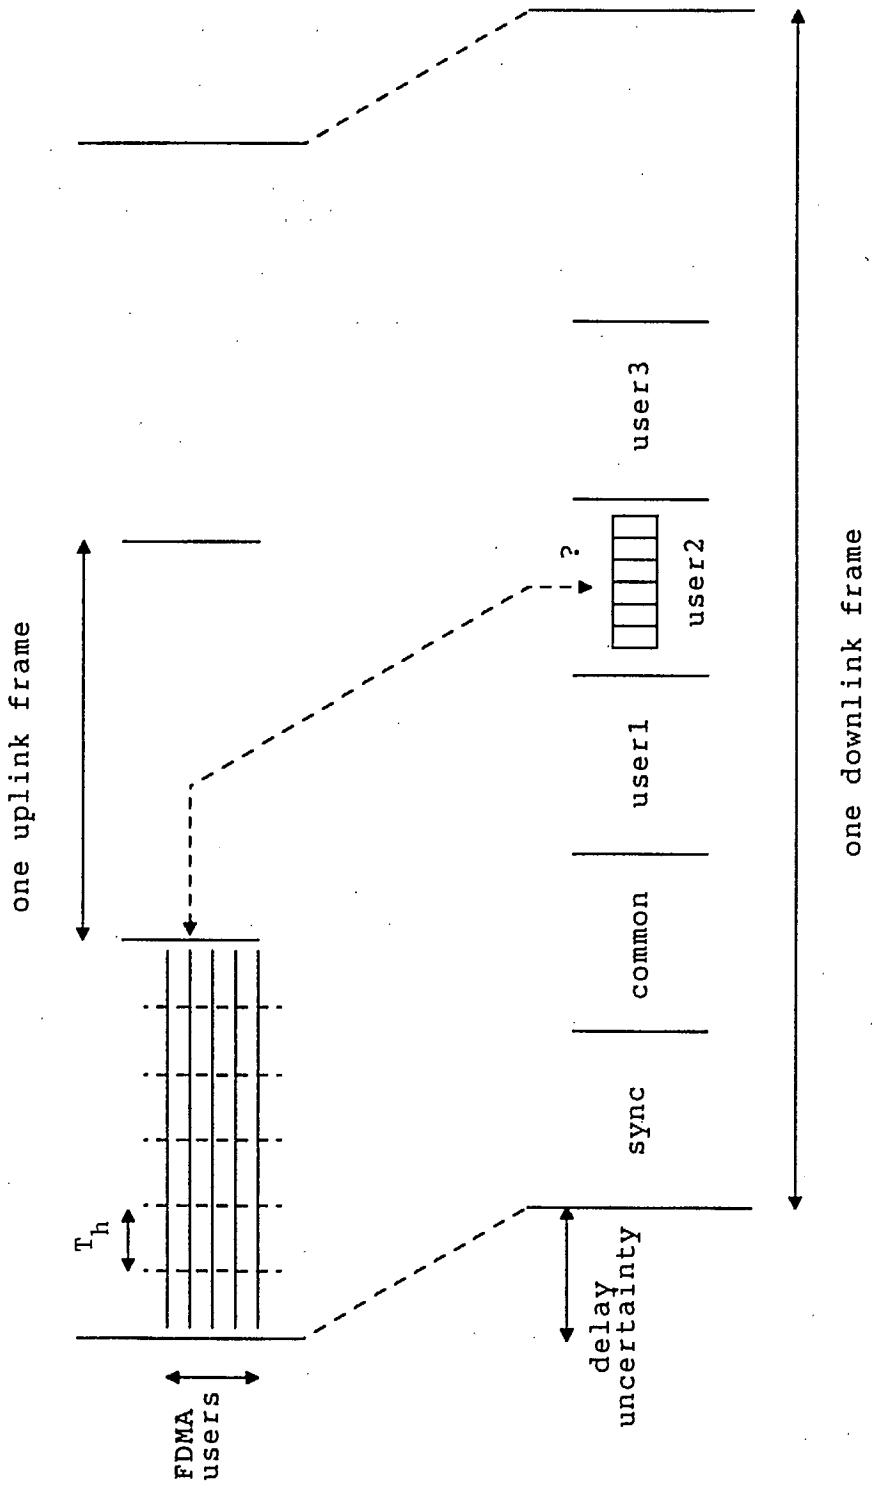

Figure 9: Initial correspondence problem.

other places in the downlink data where the detected tone agrees with the transmitted sync tone by random chance. This is related to P_{fd}. If P_{fd} is 1/8 ($M^* = 8$) , then spending at least ten hops in a row may be sufficient. If P_{fd} is only 1/2 ($M^* = 2$), we clearly must spend a much larger number of hops at each cell. For the initial uncertainty of ± 5 hops assumed for this system, and with ΔT no smaller than $T_h/2$, we will have at most twenty cells to probe. When pipelined as shown in Figure 7, we would then spend at least $4000/20 = 200$ hops in a cell. Under these circumstances, the ± 3 hops of uncertainty from the initial correspondence problem may safely be ignored. For example, we can send 206 hops worth of signal, but examine only the "central" 200 hops of these 206. This ensures that we know the correspondence is exact. The wasted 6 hops is an insignificant overhead, and is ignored in the remainder of the report.

5.4 Choosing the other $M^* - 1$ bins

It seems quite natural to select the $M^* - 1$ tone bins with which to compare from the unused in the user channel. Of course if $M^* > M$, we will have to use empty bins set aside by the system for this purpose. Note that these may be shared by all users so that the overhead this represents is insignificant. If we do use bins from the actual user channel, we must be aware of the effect of spectral spillover on the sync results. Time alignment errors τ cause a portion of the user tone's energy to appear in adjacent tone bins. The degree of spillover depends on the windowing assumed, and has been evaluated for the Kaiser-Bessel windowing (factor 1.5) assumed throughout. The ratio between the signal energy recieved in the sync tone bin and adjacent bins has been tabulated in Table 1. Smaller errors τ generate smaller spectral spillover.

The presence of signal energy in the adjacent tone bins implies that the probability of correctly detecting the user sync tone P_d will be reduced (there is less chance that it will be the best of the M^* bins). Whether this hinders or helps the sync acquisition depends on the size of the search step size ΔT. For step sizes $\Delta T = T_h/2$, we must encounter one cell that contains at least 75% of the tone; adjacent cells will contain no more than 25% of the tone. Spillover into adjacent bins is now beneficial in that it lowers P_d (relatively

speaking) much less for the "correct" cell (that causes little spillover) than it does for the adjacent ("incorrect") cells (that experience higher spillover). This lowers the probability of false detection in the incorrect cells while only minimally affecting P_d in the correct cell. If $\Delta T = T_h$, however, we may get a situation in which two adjacent cells each contain close to 50% of the signal. In this situation, the lowered P_d due to spillover will be significant enough to increase the probability of falsely detecting sync in the "incorrect" cell (the one with just under 50% of the signal).

We have chosen to avoid any extra complication by making the assumption that completely empty bins are provided by the system (say at the band edge). If this is not the case, then the results presented later will still be conservative as long as $\Delta T \leq T_h/2$.

Our choice of the other $M^* - 1$ bins also affects the results depending on the type of jamming assumed. Discussion on this is deferred until the different types of jammers are defined in the next section.

6 Evaluation of Sync Strategies

In this section, we present the results on the mean and distribution of sync acquisition times for the sync strategies that were defined during the study. We also show results on distribution of residual timing errors that represent the starting point for the subsequent fine-sync procedure (as in [5], [6]). Also included is discussion aimed at clarifying the underlying reasons for the results shown in the curves, and discussion to explain the choices made in defining the sync procedures.

For all the following results, we take $M^* = M$.

6.1 Jamming and Noise Definitions, and finding P_d

We have considered both partial band noise and tone jammers in the presence of system background noise. We follow the definitions provided in [4]. The received signal tone is

given by

$$s(t) = \sqrt{2}s\cos(2\pi f_s t + \phi) \tag{5}$$

over the hop duration T_h, giving a received signal energy per hop of $E_h = s^2 T_h$. Phase ϕ is random from hop to hop.

For partial-band noise jamming, the jammer fills a contiguous fraction γ of the hop band. For multi-tone jamming, different distributions of the jamming tones are are possible. We consider the two worst types of tone jamming as defined in [4]. The total number of tone bins is given by $N_b = W_T/W_{bin}$, where W_T is the total width of the hop band, $W_T = f_{max} - f_{min}$, and W_{bin} is the spacing of the user tones. We assume contiguous allocations of user tone bins such that user channels of width MW_{bin} are defined. A Type I jammer places his tones randomly in the bins such that some fraction γ of the total N_b bins contain a tone; $\gamma = 1$ corresponds to full-band jamming. A Type III jammer differs in ensuring that only one jamming tone appears in any user channel that is jammed. This results in a fraction β of the user channels being jammed, where $\beta = M\gamma$; $\beta = 1$ corresponds to full-band jamming.

The received signal to noise ratio due to background system noise, SNR, is defined as E_h/N_o where N_o is the two-sided noise spectral density. The signal to jamming ratio for partial band noise jamming is defined as E_h/J_o where J_o is the equivalent noise spectral density that would result if the jammer spread his total power uniformly over the entire hop band W_T; the SJR on a jammed hop is just γ times this. For tone jamming, SJR is defined by

$$SJR_{MT} = s^2/(\gamma a^2) \tag{6}$$

where a^2 is the power in the jamming tone that falls in the user tone bin. For type I MT-jamming, an SJR of 0 dB with $\gamma = 1$ corresponds to jamming tones of the same power (and amplitude) as the user signal tone. For type III MT-jamming, an SJR of 0 dB with $\beta = 1$ corresponds to jamming tones M times the power of the user tone. Type III jamming therefore has an inherent jamming advantage over type I by a factor of M.

In our best 1-of-M^* detection strategy, the M^* bins may or may not be contiguous. If they are not, it is possible with type III MT-jamming (as defined in [4]) to end up with more than one of the M^* bins jammed (looking like type I, but with a jammer advantage

in power by a factor M over type I). To avoid this additional complication, we will redefine type III MT-jamming to mean that no more than one of our M^* bins are jammed, and type I to mean one or more of these bins can be simultaneously jammed.

We use the analysis of [4] to compute the probability of correctly detecting the sync tone on any one hop, P_d. To do this, we need to account for the fact that the signal is only present for a portion $T_h(1-\tau)$ of the hop, whereas [4] implicitly assumes $\tau = 0$. We also adjust for the fact that we assume Kaiser-Bessel windowing, whereas [4] assumes rectangular windowing. We always assume that the jamming tone, if present at all in the the user's tone bin, is present for the *full* duration of the hop (a worst-case assumption). The end result is that for all types of jamming, with or without system background noise, we simply apply the energy loss tabulated earlier in Table 1 (adjusted to be 0.0 dB at $\tau = 0$ for the sync tone bin) to find the necessary E_h for the formulae of [4]. In effect, this approach is equivalent to replacing the real signal tone, which is only present for $T_h(1-\tau)$ of the total T_h hop duration, with a *full - T_h duration* signal whose height has been reduced to account for the relative energy loss of Table 1. For the cases of PBN jamming, or type III MT-jamming without background noise, this is clearly an exact approach. For type I MT-jamming, or type III MT-jamming with background noise, however, it is not intuitively obvious that this is the correct thing to do. We have, however, verified by computer computation that this simple energy scaling is in fact numerically exact for *all* these cases, although the analytical proof (which would be complicated) has not yet been attempted.

As a check on the simulation programs, we have verified that we can reproduce the probability-of-symbol-error curves that are presented in [4] for the case of no diversity. For comparison, Figures 10- 12 show the resulting probability of symbol error for *perfectly synchronized* users with $M = 2$, and for the three types of multi-tone jamming. Type II jamming is the least detrimental, and is expected to remain so as far as synchronization is concerned. We therefore do not consider it further. For all the results that follow, if system background noise is included, it is done so at an SNR of 13.35 dB, the same value used in [4]. This SNR will give a raw error symbol rate of 10^{-5} for binary FSK ($M = 2$) in the absence of jamming, and with perfect synchronization.

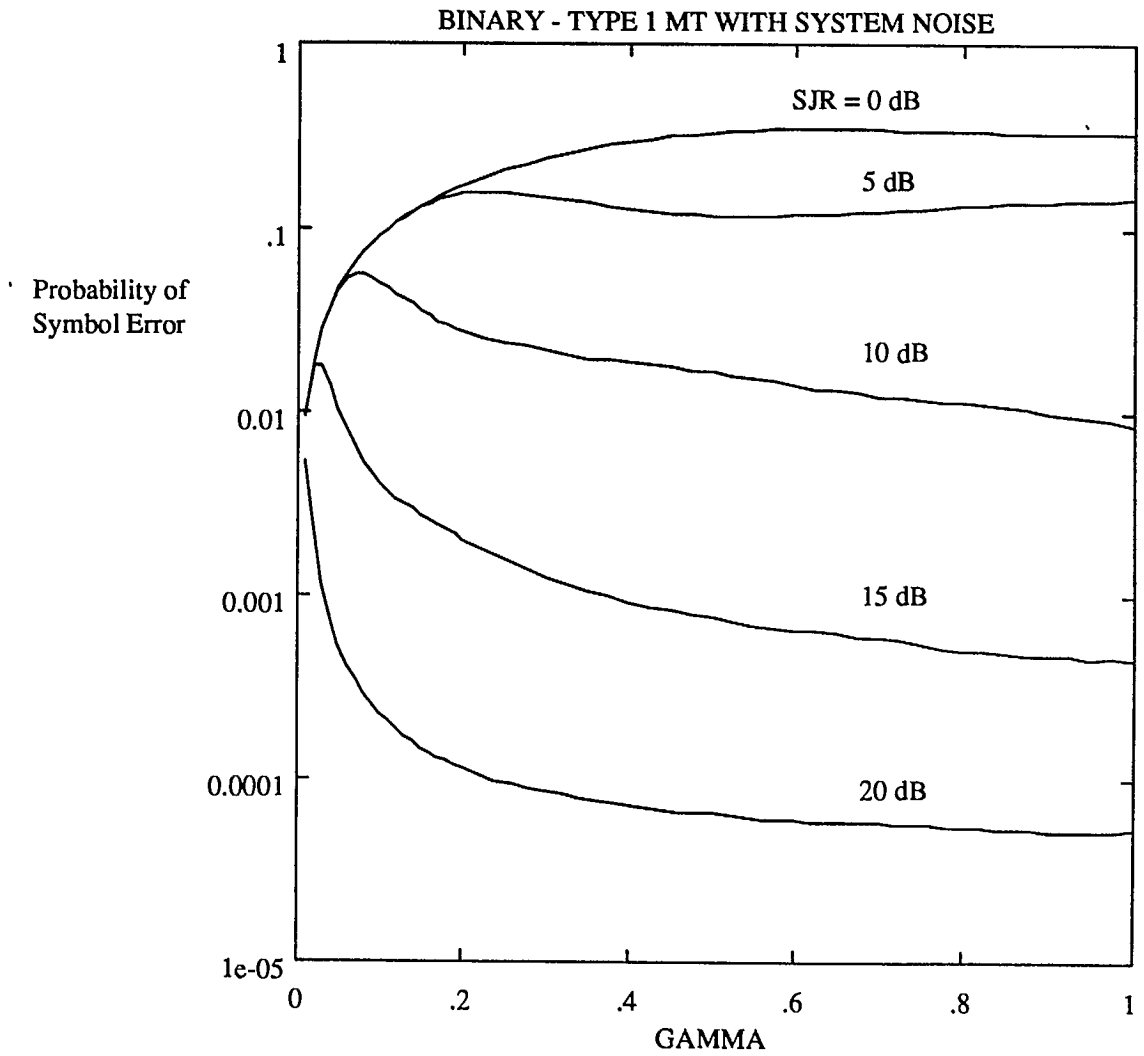

Figure 10: Symbol error probability for synchronized users, type I, M=2.

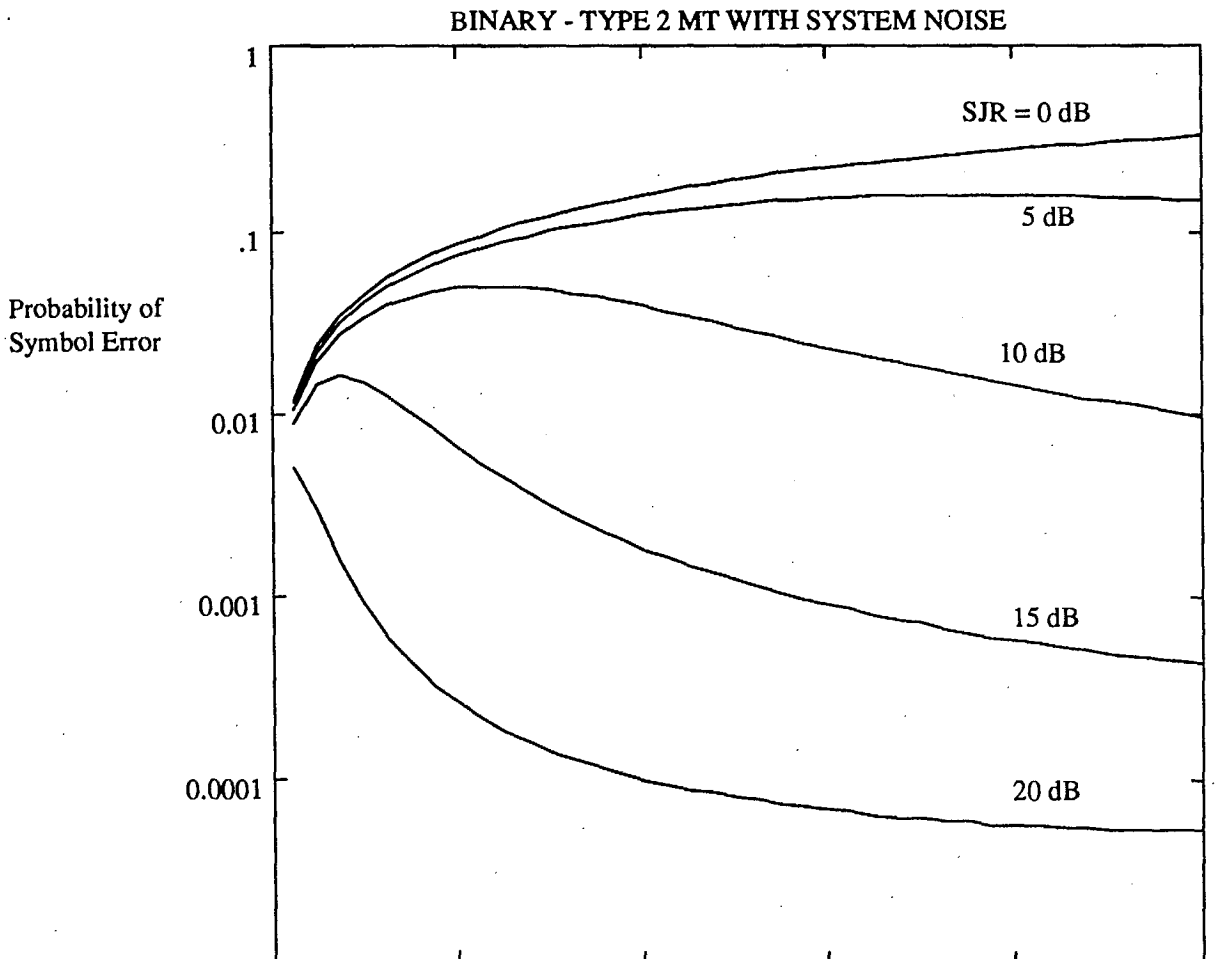

Figure 11: Symbol error probability for synchronized users, type II, M=2.

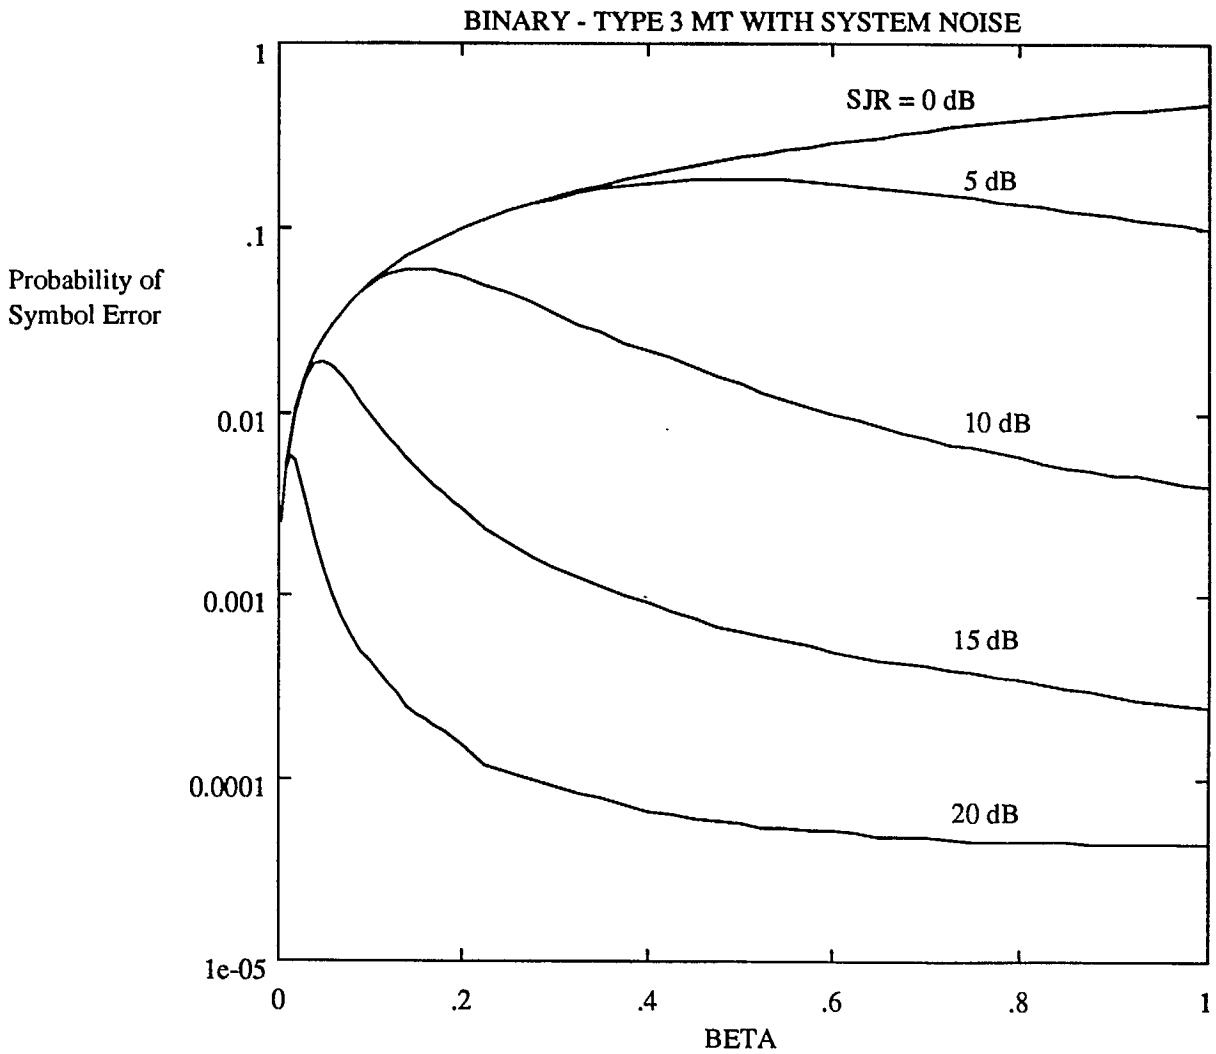

Figure 12: Symbol error probability for synchronized users, type III, M=2.

6.2 Sync Strategy 1

The first and simplest sync strategy simply visits each of the search cells in turn, stepping the hop sequence phase by $\Delta T = T_h$ or $\Delta T = T_h/2$ between cells. The number of cells visited is either 11 or 21, depending on whether ΔT is T_h or $T_h/2$. This covers the range of ± 5 hops of initial uncertainty we have assumed. At each cell the user transmits 400 hops worth of sync tone (which fills the pipeline of Figure 7 for $\Delta T = T_h$), and monitors the downlink loopback return after the propagation delay. The number of times that the sync tone bin has the highest energy of all the M^* bins is totalled, and compared to a count threshold Γ. If any of cells exhibits a count in excess of the threshold, sync is immediately declared. If none of the cells exceed the threshold, we visit them each in turn again, repeating this process until eventually one of the cells does exceed the threshold. The threshold is selected to keep the probability of false sync detection low ($\preceq 10^{-5}$).

As a starting point, we initially ignore all but the cell of best alignment, assuming that immediately adjacent cells do not contain any signal energy. Because $P_{fd} = 1/M^*$ on any one hop, we can predict what the probability of false sync detection will be after one or more visits through the cells. Because the number of visits is random, the overall probability of false acquisition, P_{facq}, will be higher for lower SJR, and will also vary with γ. Once we have the probability of correctly detecting the sync tone, P_d, and knowing the probability of falsely detecting the sync tone when not aligned is $P_{fd} = 1/M^*$, we can compute the probability of declaring synchronization in a given search cell (averaged over all $0 < \tau < 0.5$), and then compute the average time-to-acquisition. An expression for the average time to acquisition, $\overline{T_{acq}}$, considering the possibility of multiple passes can be shown to be

$$\overline{T_{acq}} = \overline{T} + \frac{t_{prop}}{P_{sync}}$$

where t_{prop} is the round trip propagation delay, $\overline{T}$ is the average time to sync given a successful first pass ($\overline{T}$ is roughly equal to $t_{prop}/2$ given the fully "pipelined" scheme; and is exactly 0.70 t_{prop} for our implementation with $\Delta T = T_h/2$), and P_{sync} is the probability of declaring sync after the N hops observation in the correct cell (i.e. averaged over all possible τ between 0 and 0.5).

Figure 13 show results for $M^* = M = 2$ with background noise and type I MT-jamming at SJR = 0.0 dB. Here we have averaged over all possible alignments for the best cell ($0 \leq \tau \leq 0.50$). Recall that $\gamma = 1$ for a type I MT-jammer at SJR = 0.0 dB corresponds to jamming tones the same height as the user sync tone. This is clearly a high jamming level. Here a threshold $\Gamma = 255$ for $N = 400$ has been found by trial and error to achieve a $P_{facq} \preceq 10^{-5}$.

At low fractions of the band jammed, γ, a large percentage of the hops are clear, and the unjammed sync tone is readily detected with the given level of system background noise. The count threshold is easily exceeded the first time the correct cell is visited, resulting in a minimum $\overline{T_{acq}}$ (which is higher for $\Delta T = T_h/2$ because of the larger number of cells in the first pass). As γ increases, the jamming tones must fall in amplitude, but more of the hops are jammed. Since the jammer tones are still bigger than the user tone, this higher γ decreases P_d, lengthening $\overline{T_{acq}}$. At $\gamma = 1$, the jammer tone has the same energy as the user tone, and all tone bins are jammed. In this case, the jammer tone in the sync tone bin adds to the signal energy, *improving* P_d, and lowering $\overline{T_{acq}}$. Optimum jamming occurs at a somewhat smaller γ, around $\gamma = 0.8$.

These curves also clearly show the beneficial effect of a smaller step size; the average time to acquisition $\overline{T_{acq}}$ is greatly reduced. The larger number of cells searched is compensated for by the fact that the worst-case fractional energy loss due to residual timing error τ in the best cell is much less with the smaller step size. Figure 14 shows probability of false acquisition on an empty cell, P^e_{facq}, corresponding to the $\overline{T_{acq}}$ of Figure 13. Note that P_{facq} remains below the target of 10^{-5} for all γ when the smaller step size is used. For SJR $\succeq$ 5.0 dB, all the curves in Figures 13- 14 will become virtually flat lines, indicating the absolute minimum $\overline{T_{acq}}$ has been achieved. At this higher SJR, the jammer tones are smaller than the user tone for $\gamma \geq .316$. At higher γ then, the jammer tones are too small to result in a significant number of missed sync tone detections; and at lower γ, enough of the hops are clear from jamming so that the sync tone presence is readily detected. The minimum $\overline{T_{acq}}$ is therefore achieved for all γ.

Figures 15- 16 show mean acquistion times for for type III MT-jamming. Now $\beta = 1$ at

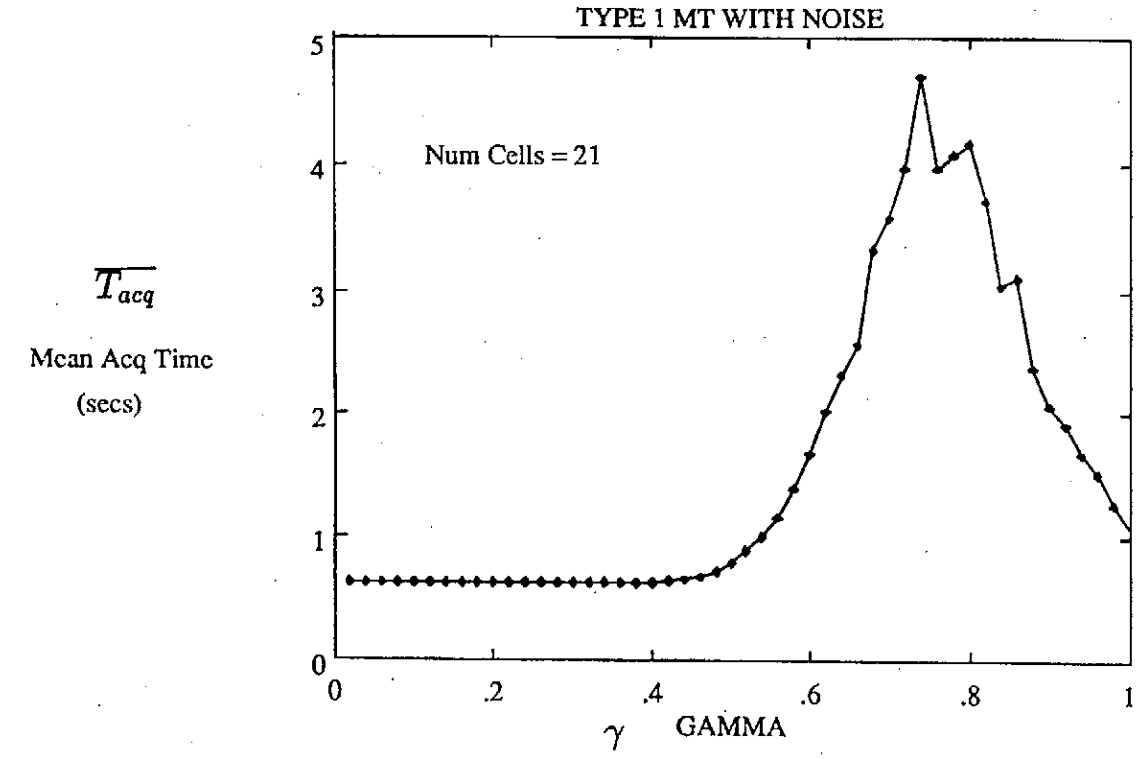

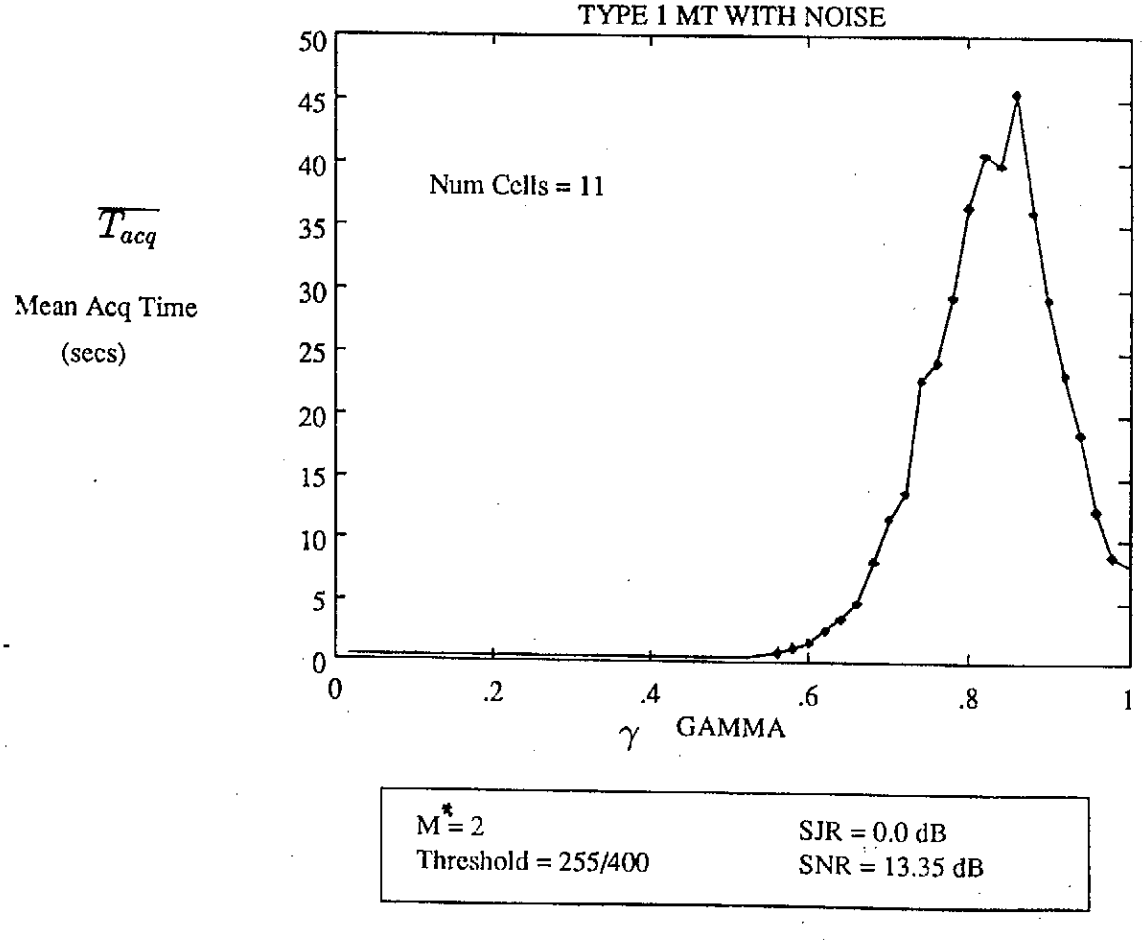

Figure 13: Mean acquisition time versus fraction of band jammed, type I.

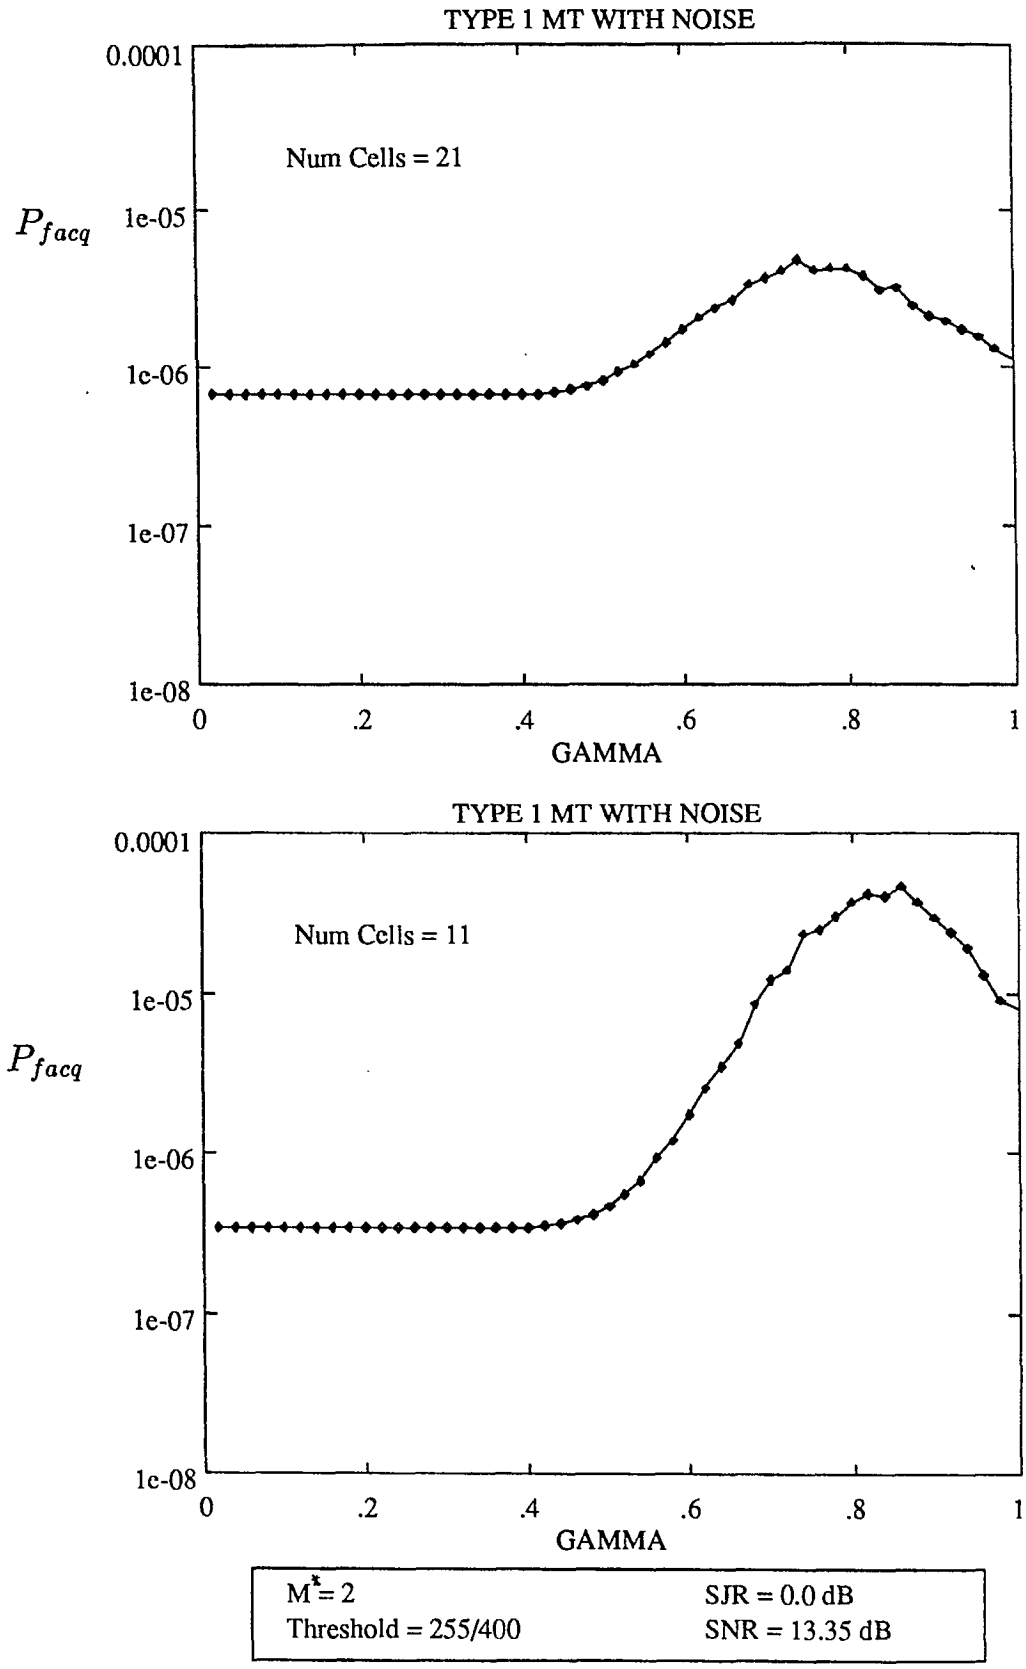

Figure 14: Probability of false acquisition, type I.

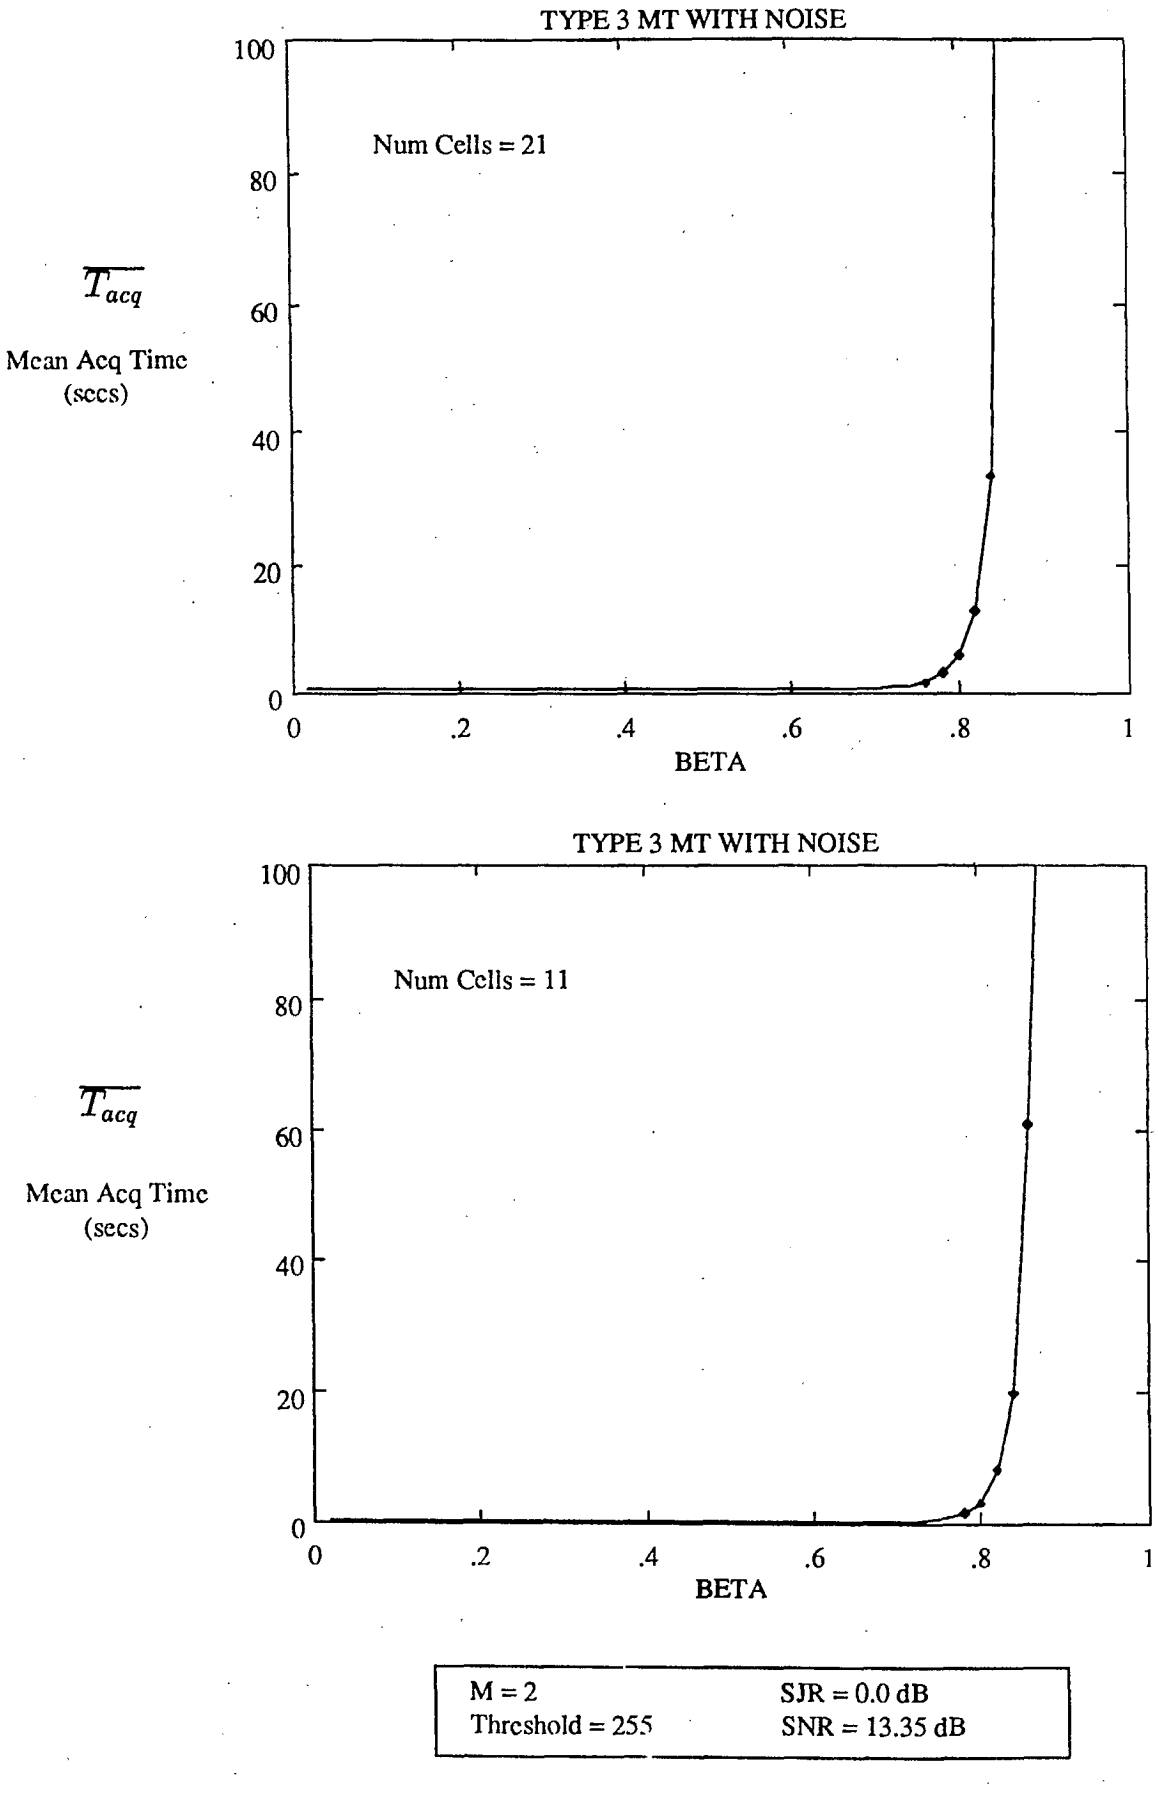

Figure 15: Mean acquisition time, type III at SJR=0.0 dB.

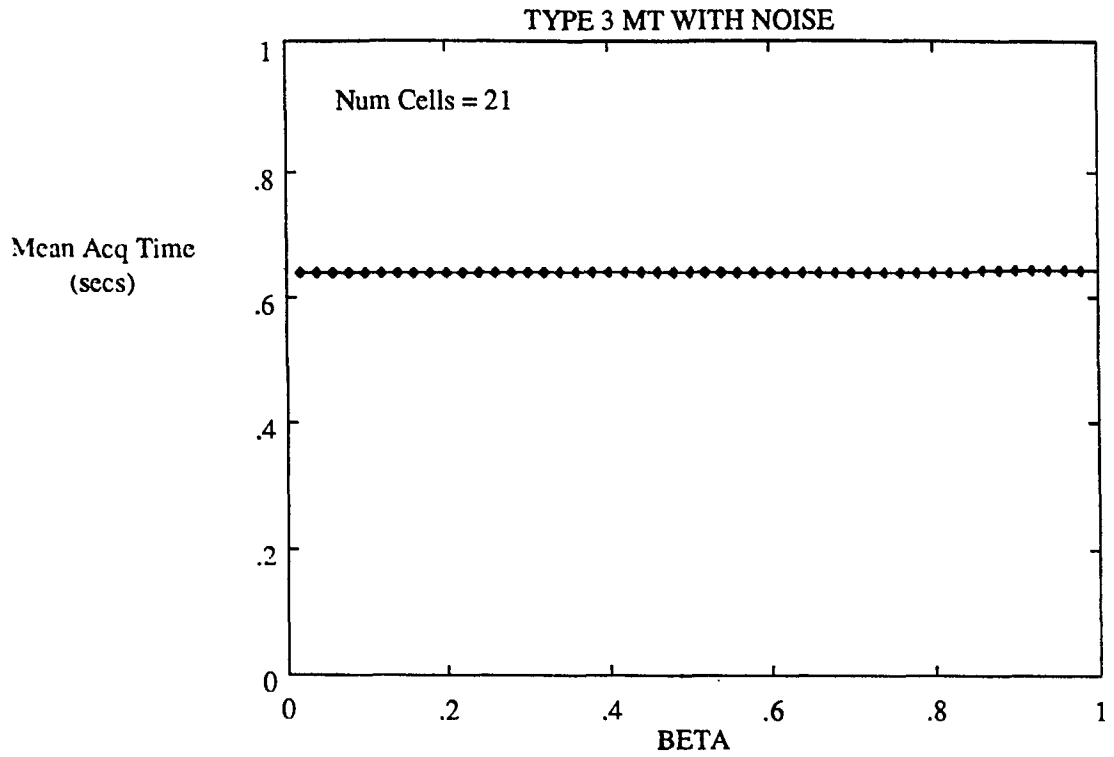

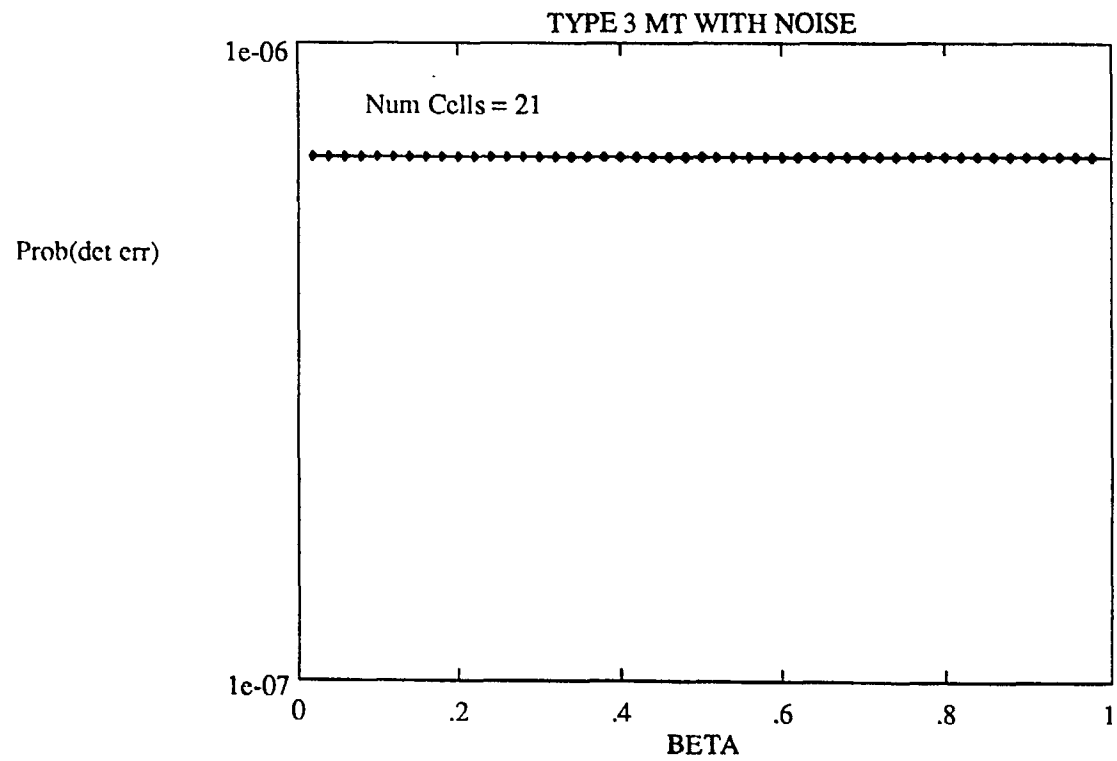

M = 2	SJR = 5 dB
Threshold = 255/400	SNR = 13.35 dB

Figure 16: Mean acquisition time, type III at SJR=5.0 dB.

SJR=0.0 dB corresponds to a jammer tone that is M times the energy of the signal tone. Again at low β, most of the hops are clear, and the sync tone is reliably detected. In fact, we only need just more than 25% of the hop band clear ($\beta \preceq 0.75$) to bias the number of detection counts to exceed the threshold. That is, at $\beta = .75$, we should expect detections on almost all of the 25% (100) clear hops, plus random detections on a fraction $1/M^* = 1/2$ of the other 75% (300) jammed hops, for a total of about 250 detections on average, virtually right at the threshold of $\Gamma = 255$. Any lower β, and the threshold is easily exceeded. At higher β, the situation is clearly hopeless because SJR = 0.0 dB has the jammer tones $\geq M$ times bigger than the user tone right up to $\beta = 1$. Note that the curves are much steeper than the type I MT-jamming case. At a lower jammer total power, SJR $\succeq$ 5.0 dB, the jammer tones at $\beta = 1$ (with $M = 2$) are now smaller than the signal tone, and have little chance of causing detection errors. At lower β, the jammer tones will be larger and cause missed detections on jammed hops, but now a sufficient fraction of the hops are clear, causing the count threshold to be easily exceeded. The result is that for SJR $\succeq$ 5 dB, for *all* β we readily achieve the minimum $\overline{T_{acq}}$ seconds for a step size of $T_h/2$, with $P_{facq} < 10^{-6}$ (again, the curves become flat lines : Figure 16).

The high worst-case average times to acquisition in the previous results indicate that the cells are being visited many times; that is, threshold $\Gamma = 255$ is not readily exceeded on any one visit. This must be so in order to keep P^e_{facq} low. By going to M^*=4 or 8 (with Γ=150 or Γ=100 for 400 hops, respectively), we achieve a beneficial tradeoff: P_d does drop somewhat, but $P_{fd} = 1/M^*$ is lowered even more. With these thresholds, the first visit to the best cell is virtually always sufficient, even with $\Delta T = T_h$, , while P^e_{facq} is even lower (see Figure 17). Results in later sections use M^*=4 and $\Delta T = T_h$ on the first pass, and show the much lower $\overline{T_{acq}}$ values.

The previous results did not consider the effect of non-empty cells adjacent to the best cell (which contain non-zero signal energy). These cells can have a reasonably high probability of exceeding the count threshold, and may raise P_{facq} to unacceptable levels (where we consider declaration of acquisition in a cell having $\tau \geq 0.50$ to represent a false sync detection). This was initially investigated for the noise-only case at SNR=0.0 dB. For a step size of $\Delta T = T_h$,

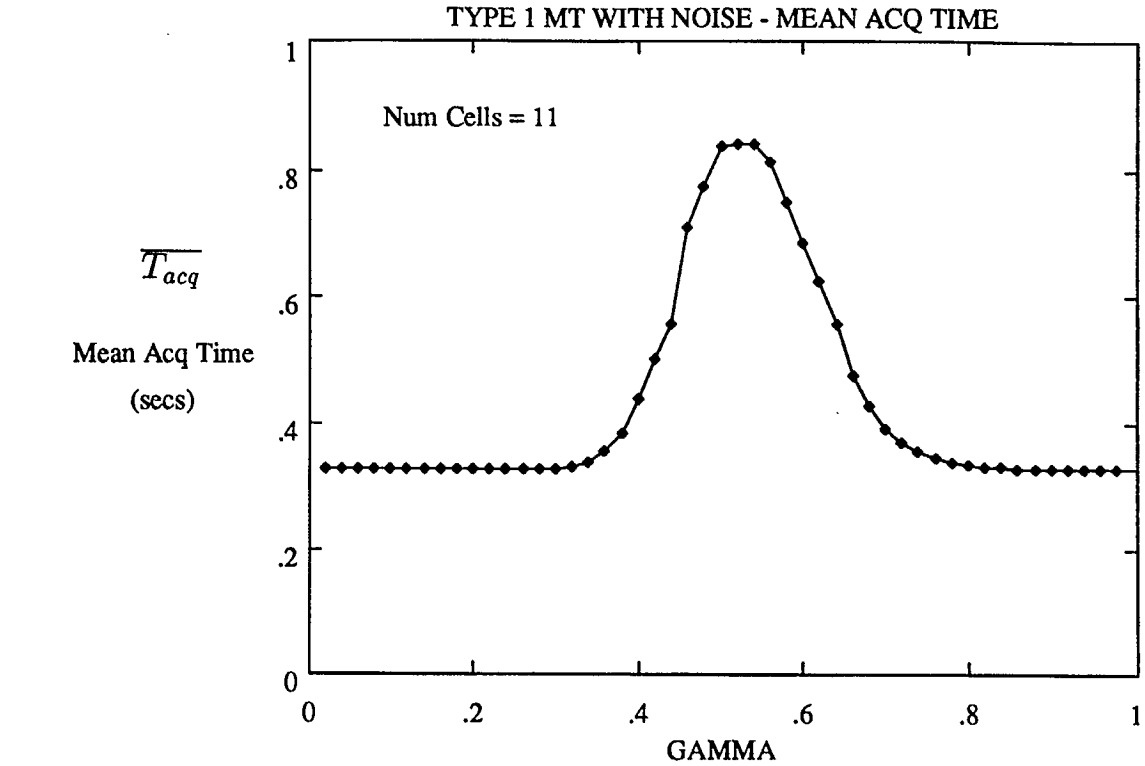

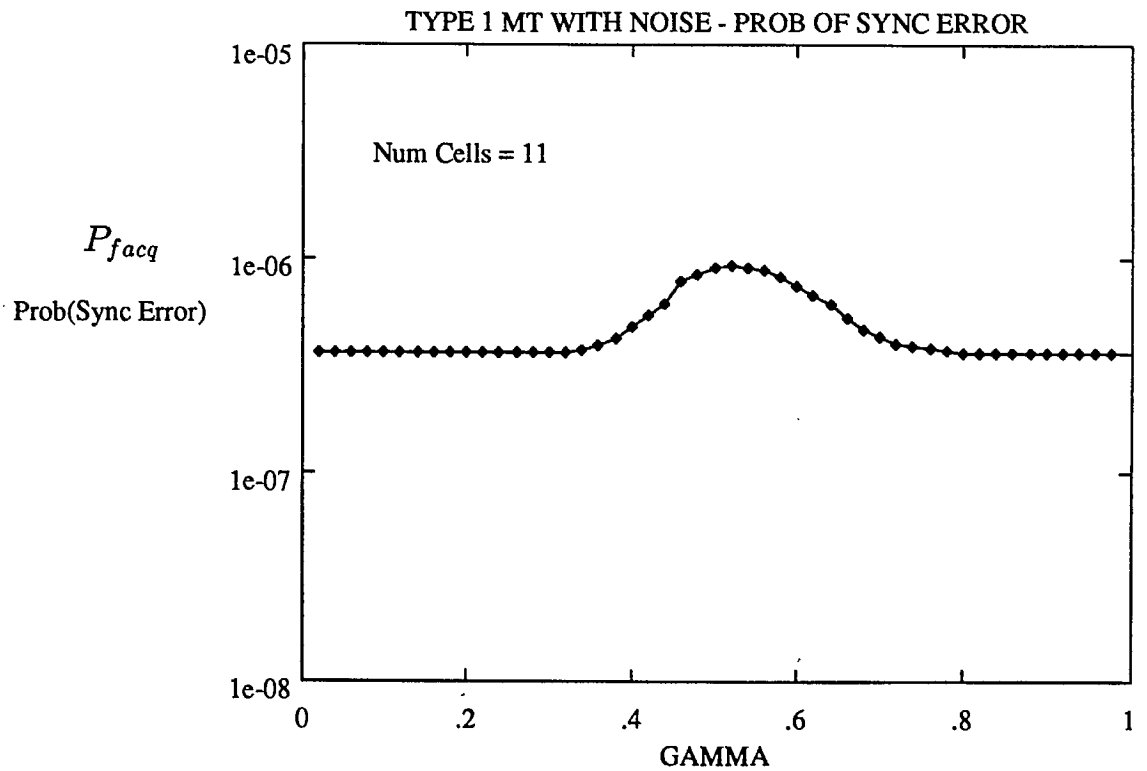

M^*= 4	SJR = 0.0 dB
Threshold = 150/400	SNR = 13.35 dB

Figure 17: Mean acquisition time and P^e_{facq}, type I, M^*=4, $\Delta T = T_h$.

worst-case alignment creates two adjacent cells, one having signal energy for a duration $T_h/2 + \epsilon$ (the "correct" cell), and the other having signal for $T_h/2 - \epsilon$ (the "incorrect" cell), for epsilon a small value. In this situation, $P_{facq} \approx 0.50$. If we go to a step size of $\Delta = T_h/2$, however, the worst case alignment is 75% in the correct cell, and 25% in two adjacent incorrect cells. Now P_{facq} drops back to around 10^{-6} for a threshold $\Gamma = 260$ out of the 400 hops.

Figure 18 shows the distribution of the resulting residual timing errors τ that will be passed to the fine-sync procedure with full-band noise jamming at SNR = 0.0 dB, $M^* = 2$. In this case, the threshold detection with $\Delta = T_h/2$ achieves $\overline{T_{acq}} = .618$ seconds. At SNR = 5.0 dB and 10.0 dB, $\overline{T_{acq}} = .448$ and .442 seconds respectively (near minimum), and the distributions are somewhat tighter (Figures 19- 20).

Figure 21 shows results for type I MT-jamming, $M^* = 4$, $\Delta = T_h$, with adjacent cells accounted for. The count threshold is now 150, 50 counts in excess of the 100 counts expected for cells containing no signal energy. Note the high probability of false sync declaration in the incorrect adjacent cell (peaking around $\gamma = 0.4$). This poor performance is in marked contrast to the full-band noise-only case just looked at. Figure 22 shows the resulting mean times to acquisition. Although acquisition times look good, the problem here is that sync is being declared too often on incorrect cells containing partial signal energy. Figure 23 shows that the highest number of missed detections occurs at $\gamma = .4$, where the jamming tones are 2.5 times the energy of the user tone. For $\gamma = 0.4$, Figure 24 then shows the probability of detecting the sync tone on any one hop, P_d, versus the normalized timing error τ in any cell having partial alignment. This latter figure shows the root of the problem: P_d is not very different in a cell of near-perfect alignment ($\tau \approx 0$) than it is in a cell having τ just in excess of $\tau = 0.5$. Clearly, this makes it very difficult to tell the correct from the incorrect cell. This P_d behaviour is readily explained.

The expected percentage of clear hops (none of the $M^* = 4$ bins jammed) at $\gamma = .4$ is $(1 - \gamma)^4 = .13$. For these clear hops, $P_d \approx 1$ for all but large τ (where background noise will become significant). The expected percentage of jammed hops at $\gamma = 0.4$ is $1 - (1 - \gamma)^4 = 0.87$, but this includes one, two, three, or four jamming tones in the $M^* = 4$

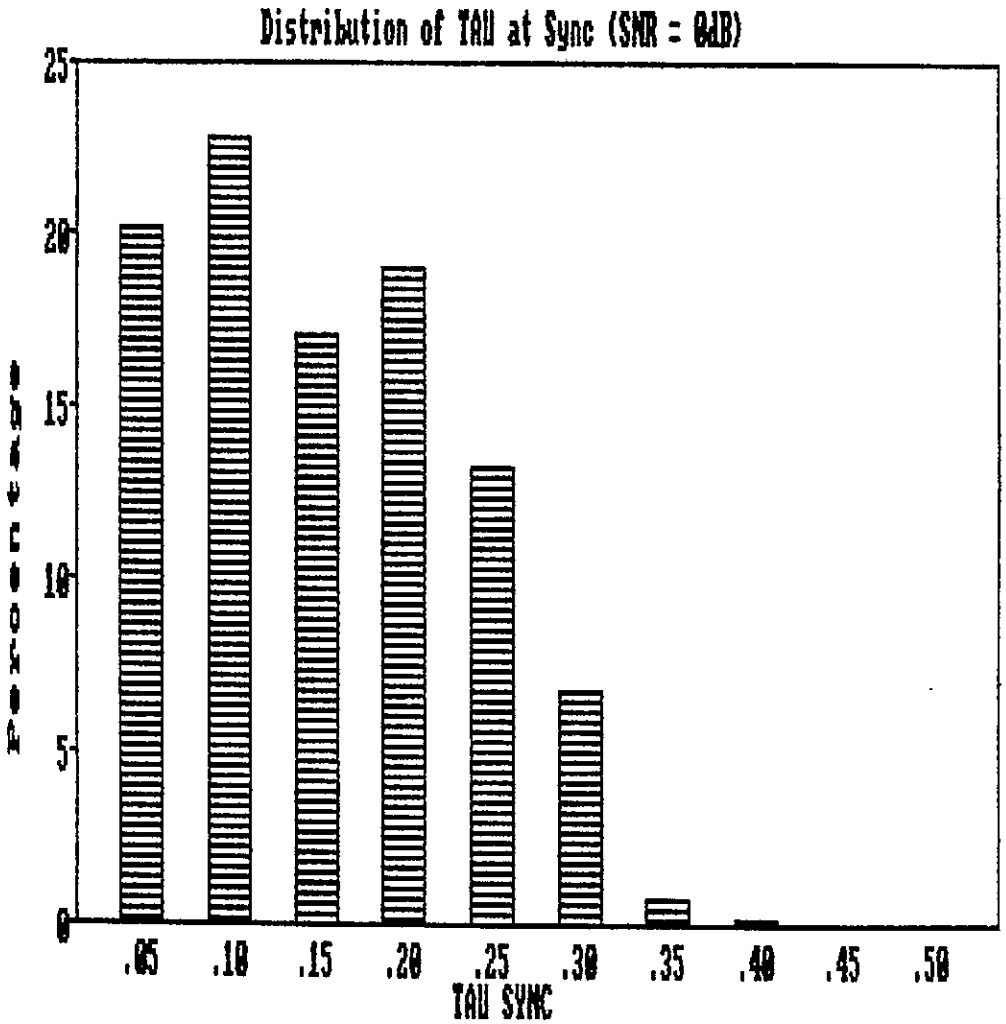

Mean Acq Time = 0.618 secs

Threshold: 255 of 400

M: 2

Figure 18: Residual timing errors with noise jamming, $\gamma = 1$, SJR= 0.0 dB

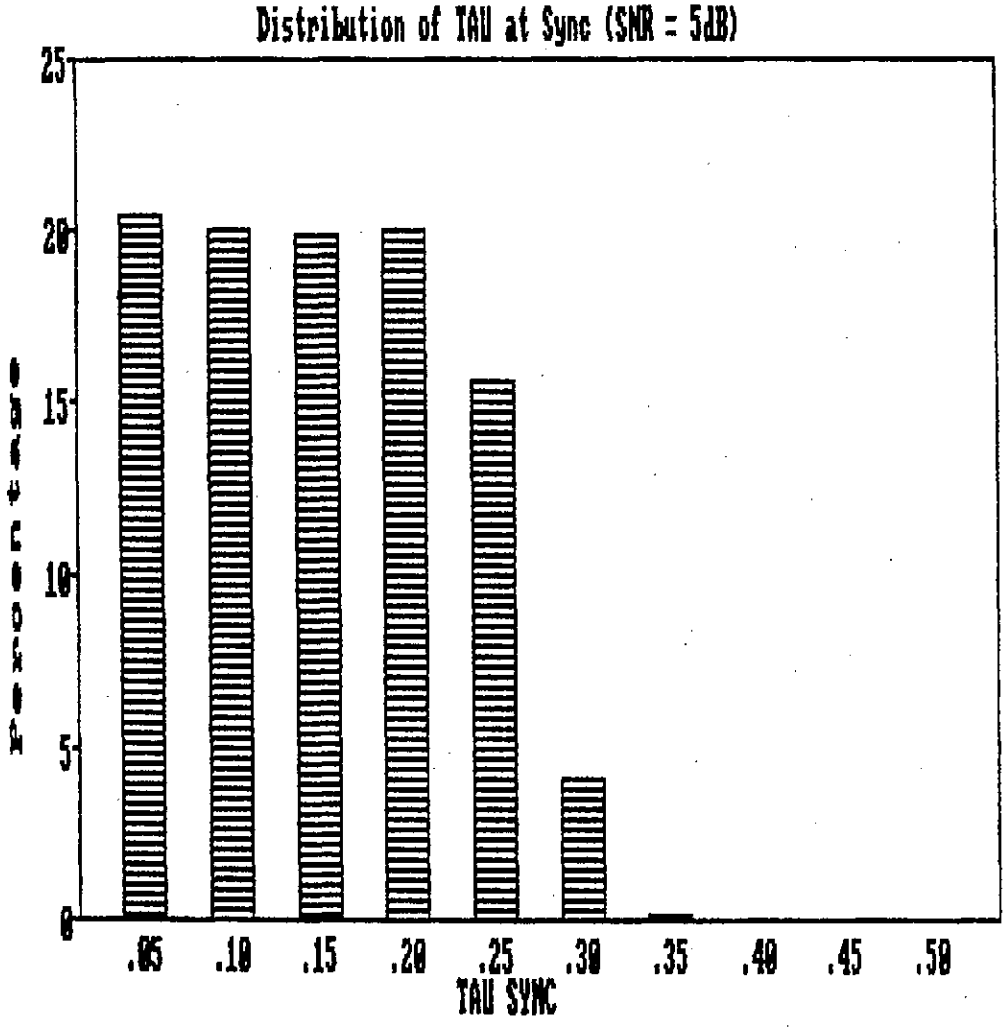

Mean Acq Time = 0.448 secs

Threshold: 255 of 400

M: 2

Figure 19: Residual timing errors with noise jamming, $\gamma = 1$, SJR= 5.0 dB

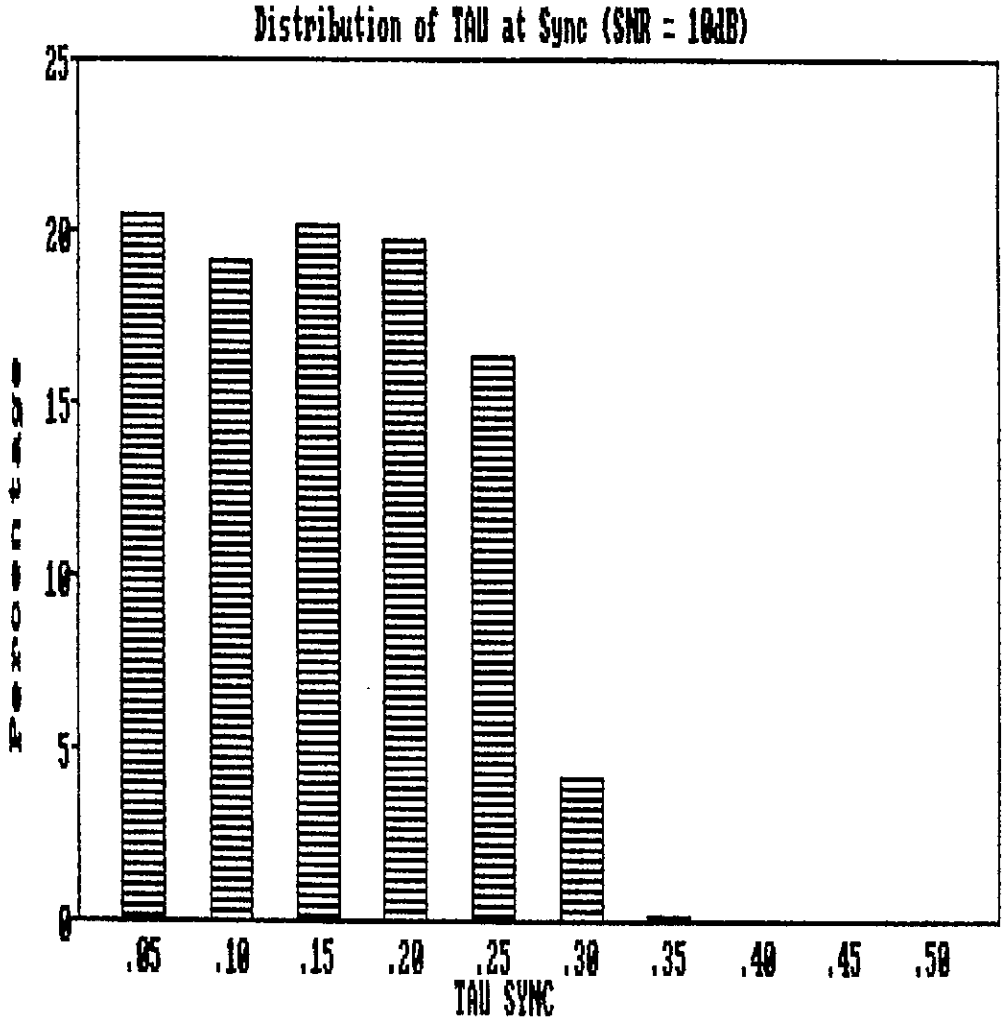

Mean Acq Time = 0.442 secs

Threshold: 255 of 400

M: 2

Figure 20: Residual timing errors with noise jamming, $\gamma = 1$, SJR= 10.0 dB

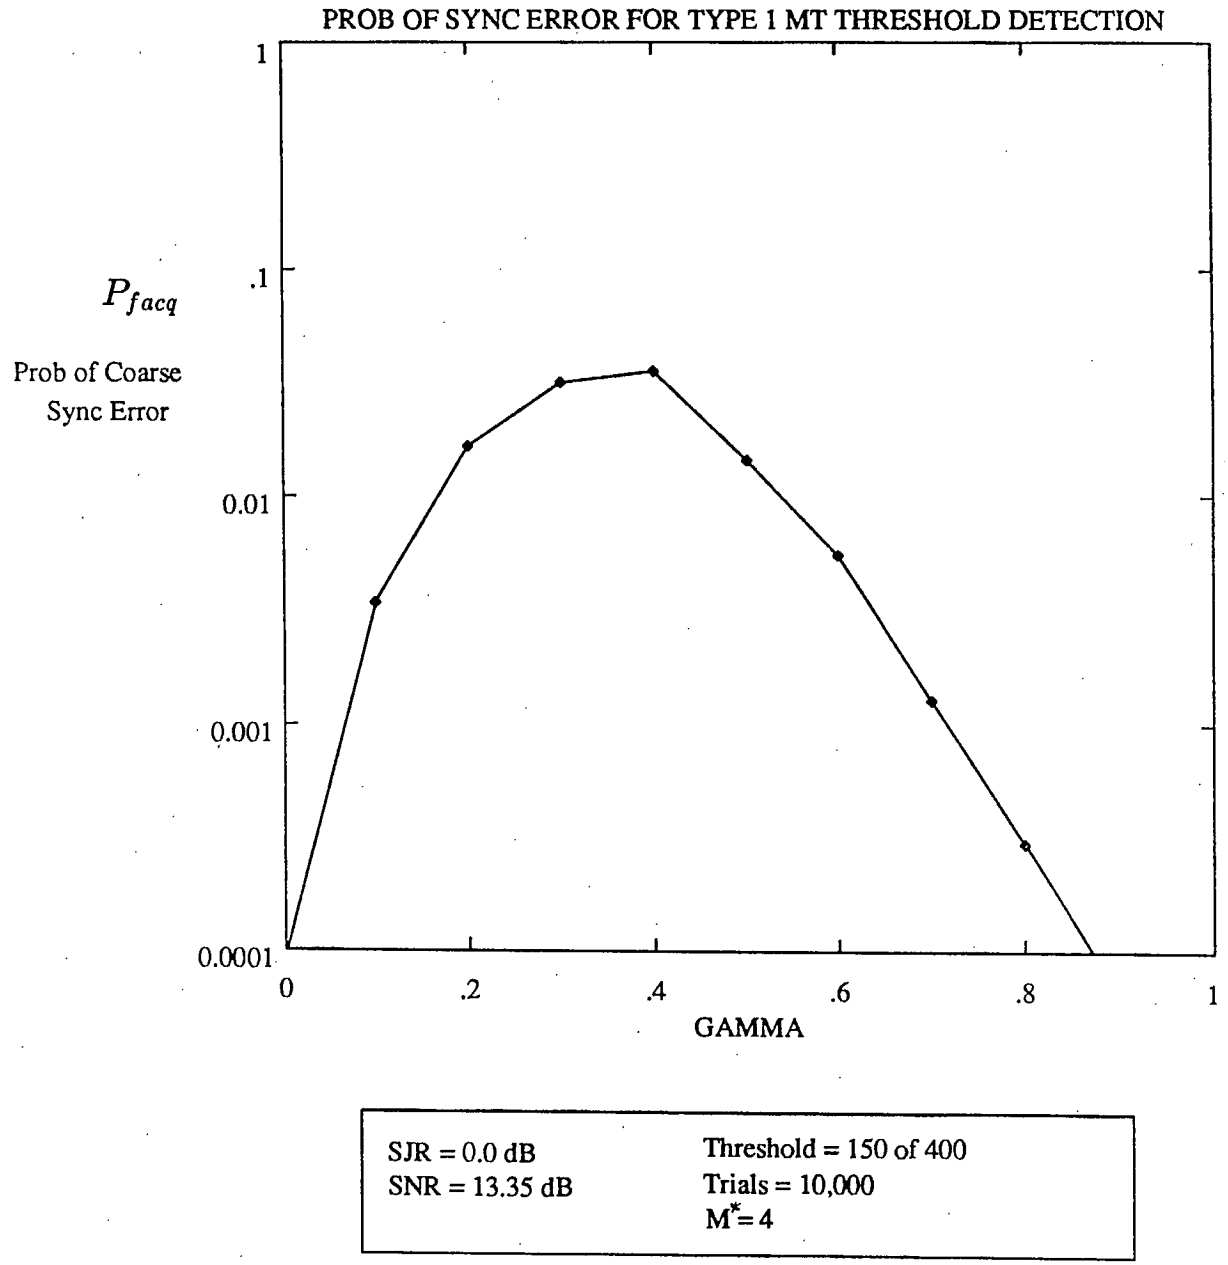

Figure 21: Probability of false acquisition, type I, $M^* = 4$, SJR= 0.0 dB

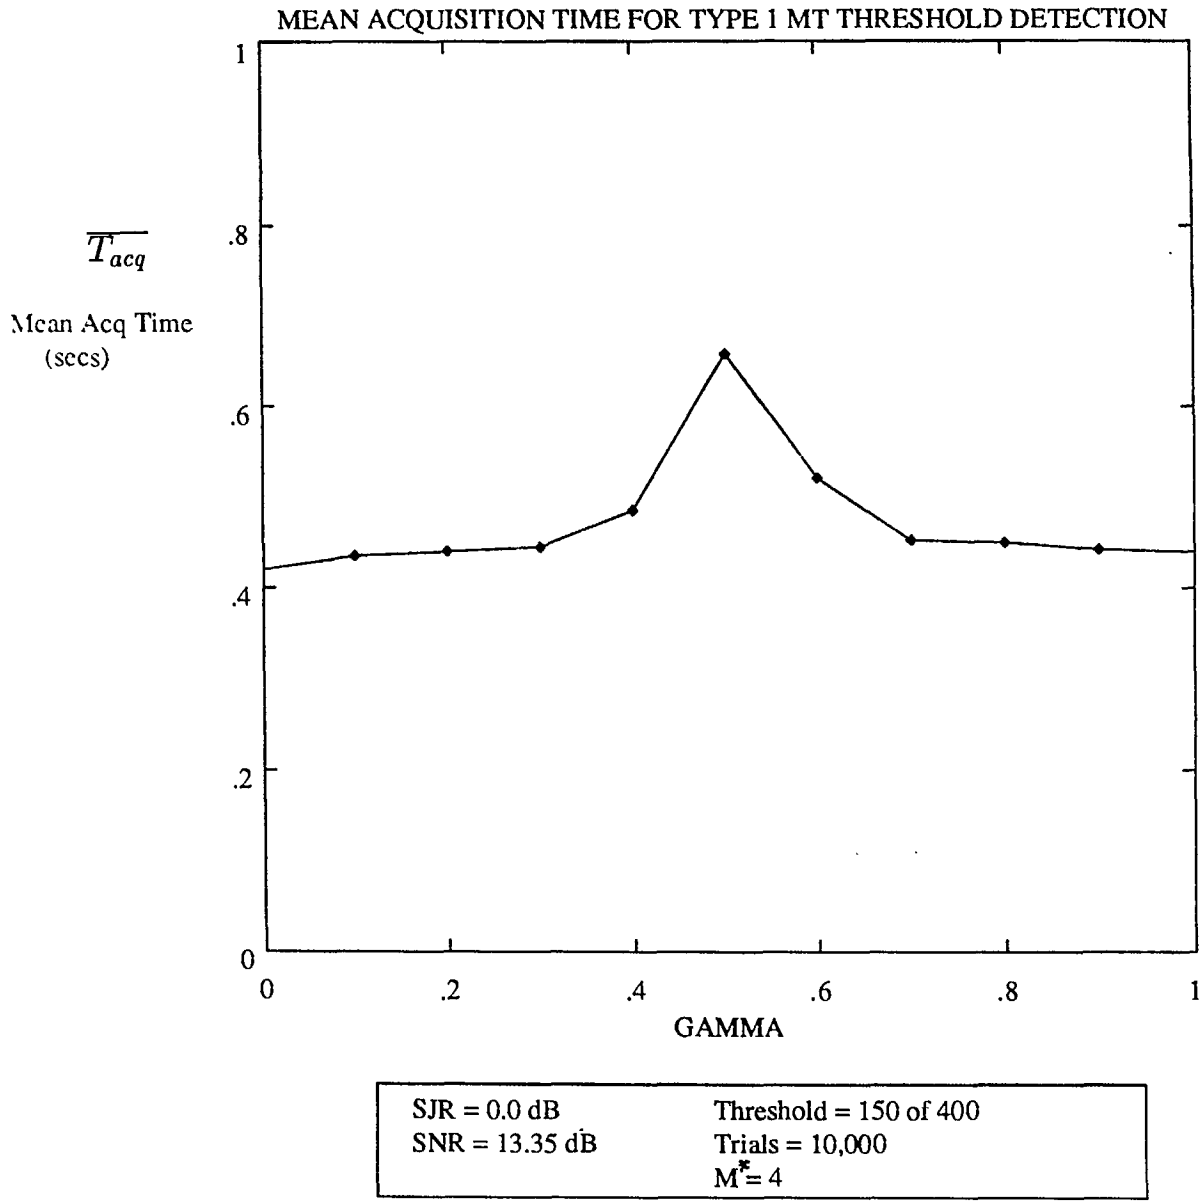

Figure 22: Mean acquisition times, type I, $M^* = 4$, SJR= 0.0 dB

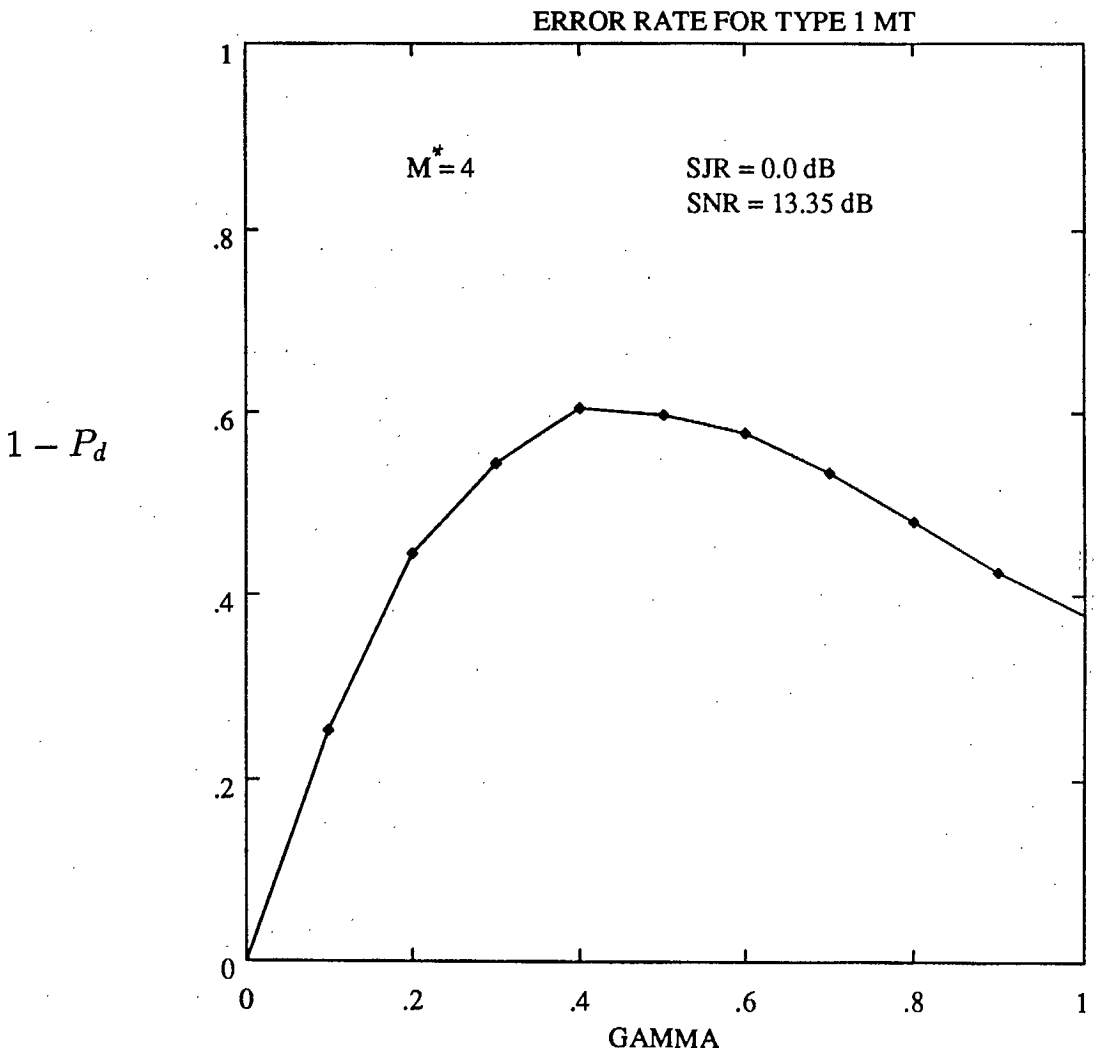

Figure 23: Probability of missed tone detection, $M^* = 4$, SJR= 0.0 dB

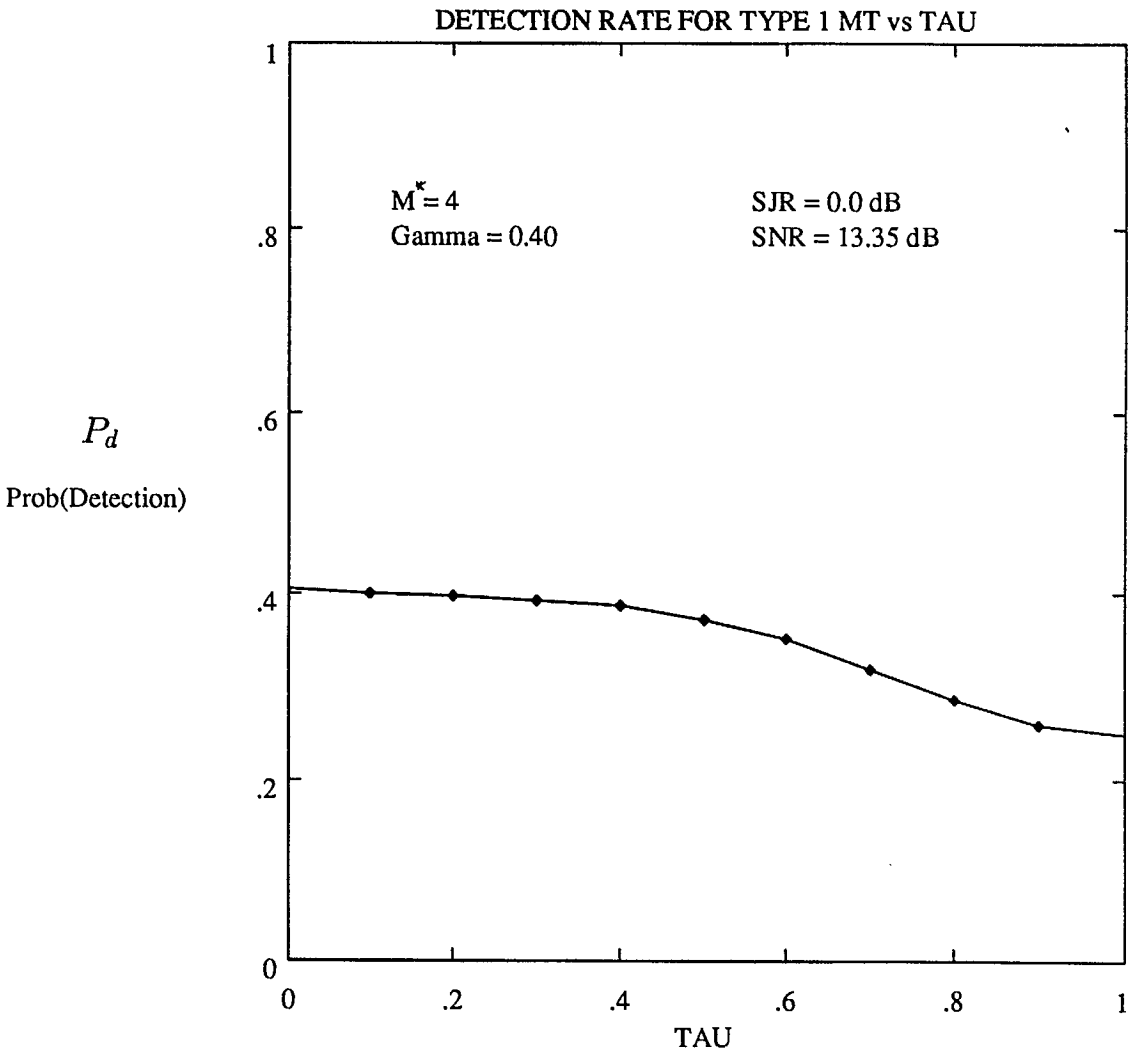

Figure 24: Probability of sync tone detection, $M^* = 4$, SJR= 0.0 dB, $\gamma = 0.40$.

bins. Since at SJR= 0.0 dB, the jammer tones are much larger than the user tone at $\gamma = 0.4$, a jammed hop having only one jamming tone in the $M^* = 4$ tone bins (total percentage of hops like this is .346) has a $P_d \approx 1/M^* = .25$ (random detections). With two, three, and four jamming tones this detection probability rises due to the higher probability that one of the jammer tones will hit the signal tone, add to it, and result in a detected energy is in excess of the other jammed bin(s). Consider the $\tau = 0$, negligible background noise case. If one of the jammer tones hits the sync bin, it will add with random phase, and result in a larger energy than any other jammer tone with probability ≈ 0.61. With two jamming tones (total fraction of hops like this is .346), there is a fifty-fifty chance of hitting the user tone bin, giving $P_d \approx 0.5(.61)$. With three jammer tones, $P_d \approx 0.75(.61)$, and with four jamming tones, $P_d \approx 1(.61)$. A simple estimate of the overall average P_d is then $P_d \approx (0.13)(1) + (.346)(.25) + (.346)(.50)(.61) + (.153)(.75)(.61) + (.026)(1)(.61) = 0.408$. In fact this turns out to be very close to the true value even with the background noise adding in, indicating that the noise effect is not large here. This probability depends only weakly on τ, decreasing as τ increases, but only slowly. The final value is $P_d = 0.25$ at $\tau = 1$ (no signal energy). This results in a low slope to the curve in Figure 24. Figure 25 shows the addition of the user tone and jammer tone in the noiseless case for $\tau = 0$ and $\tau = 0.5$. The dashed lines indicate the angle that the resultant must exceed in order that its length is less than a jammer tone in another bin. The small separation angle ϕ indicates that the probability of this event is not very different for the two τ values, causing the low P_d differential.

Figure 26 shows the $\gamma = 1$ case. Now the jamming tones are the same height as the user tones, and all tone bins are jammed. Now $P_d \approx 0.6$ at $\tau = 0$, decreasing more sharply to again end up at $P_d = 0.25$ at $\tau = 1$. As indicated in Figure 21, this removes the false-sync detection problem. At higher SJR's, P_d approaches unity at $\tau = 0$, and the slope is even steeper, again completely alleviating the problem.

Figure 27 shows the effect of different M^* on the probability of detecting the sync tone versus alignment error at $\gamma = 1.0$. Note that the P_d curve for $M^* = 2$ has a lower slope than the $M^* = 4$ curve (seen separately on Figure 26), while $M^* = 8$ has a better slope. Since

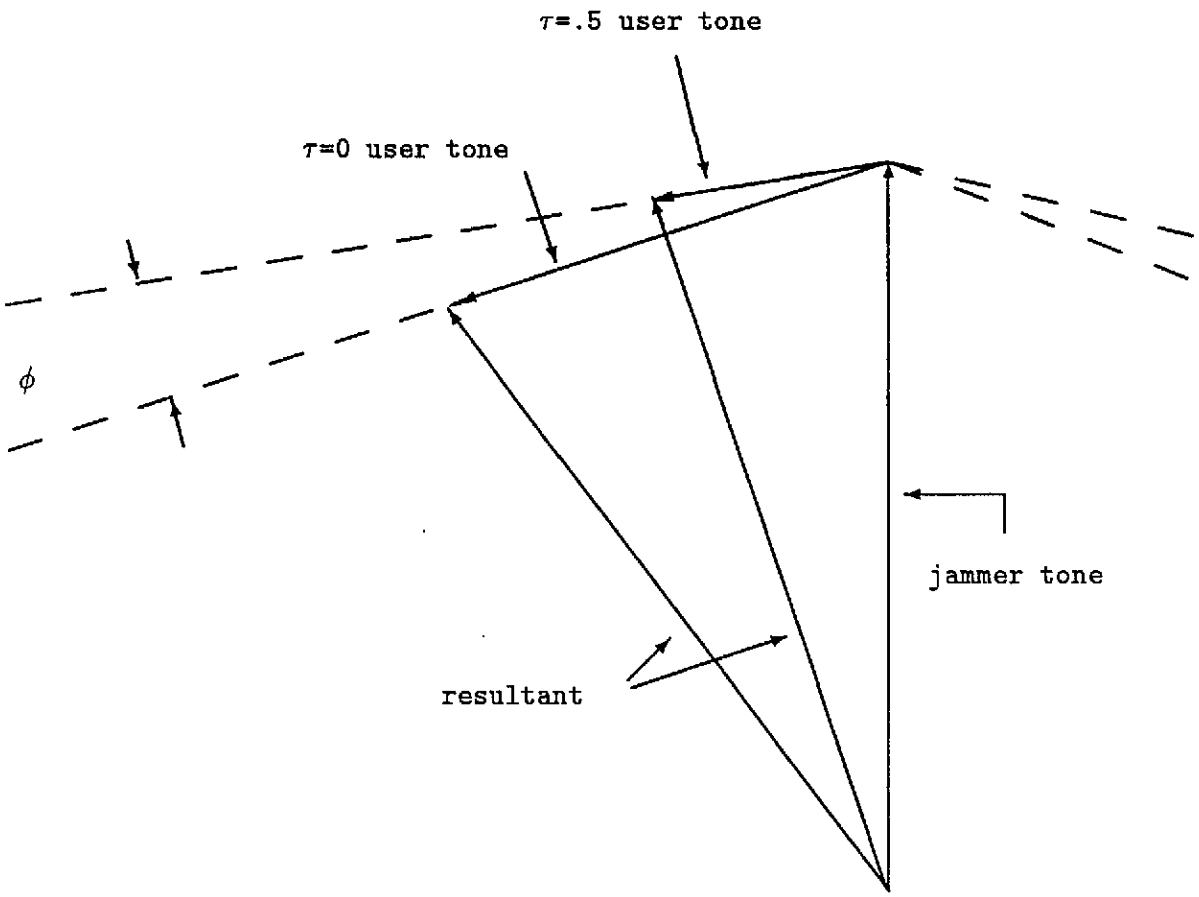

Figure 25: Addition of jammer tone and user tone, SJR= 0.0 dB, $\gamma = 0.40$.

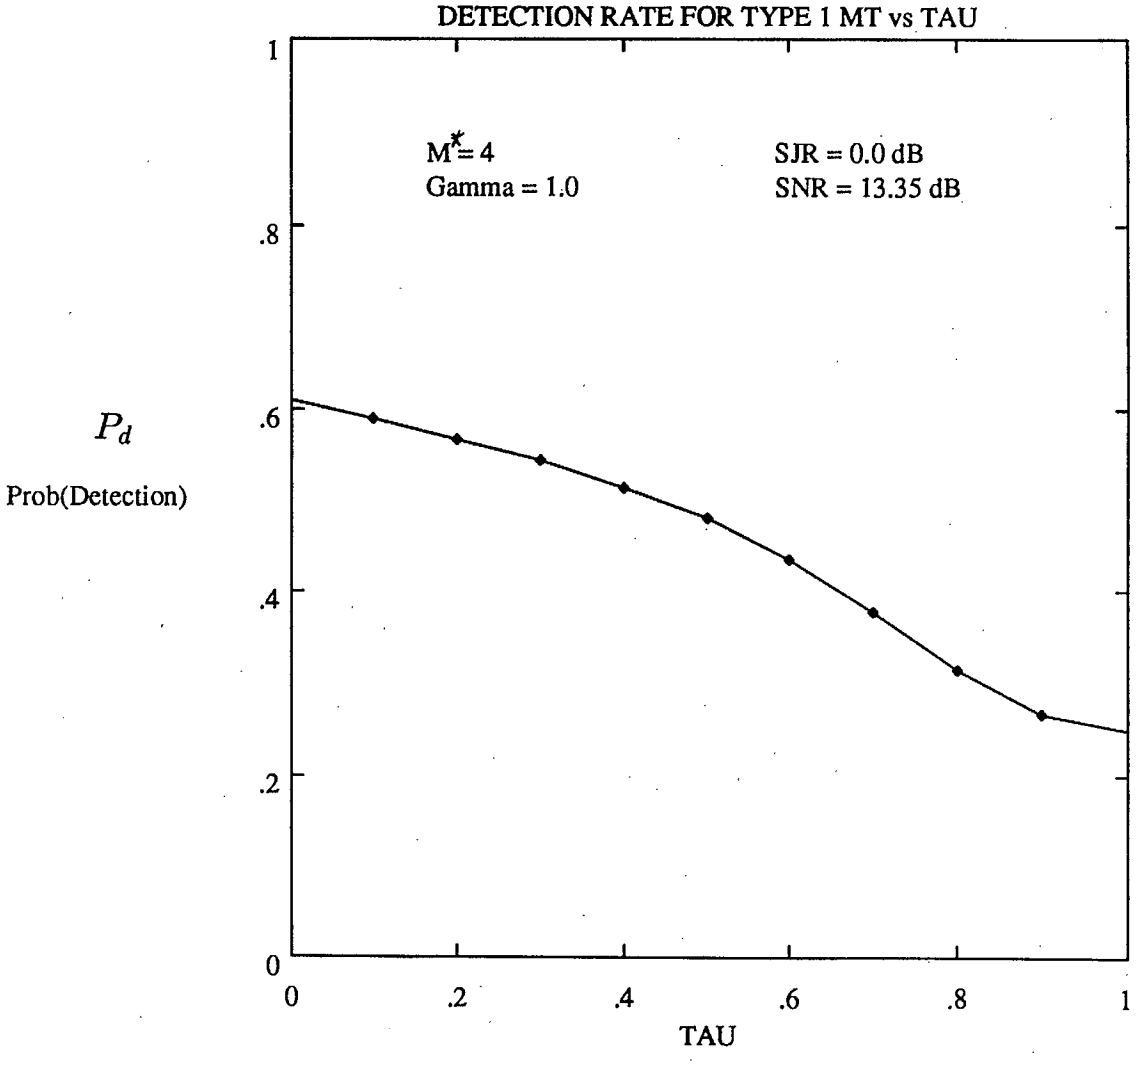

Figure 26: Probability of sync tone detection, $M^* = 4$, SJR= 0.0 dB, $\gamma = 1$.

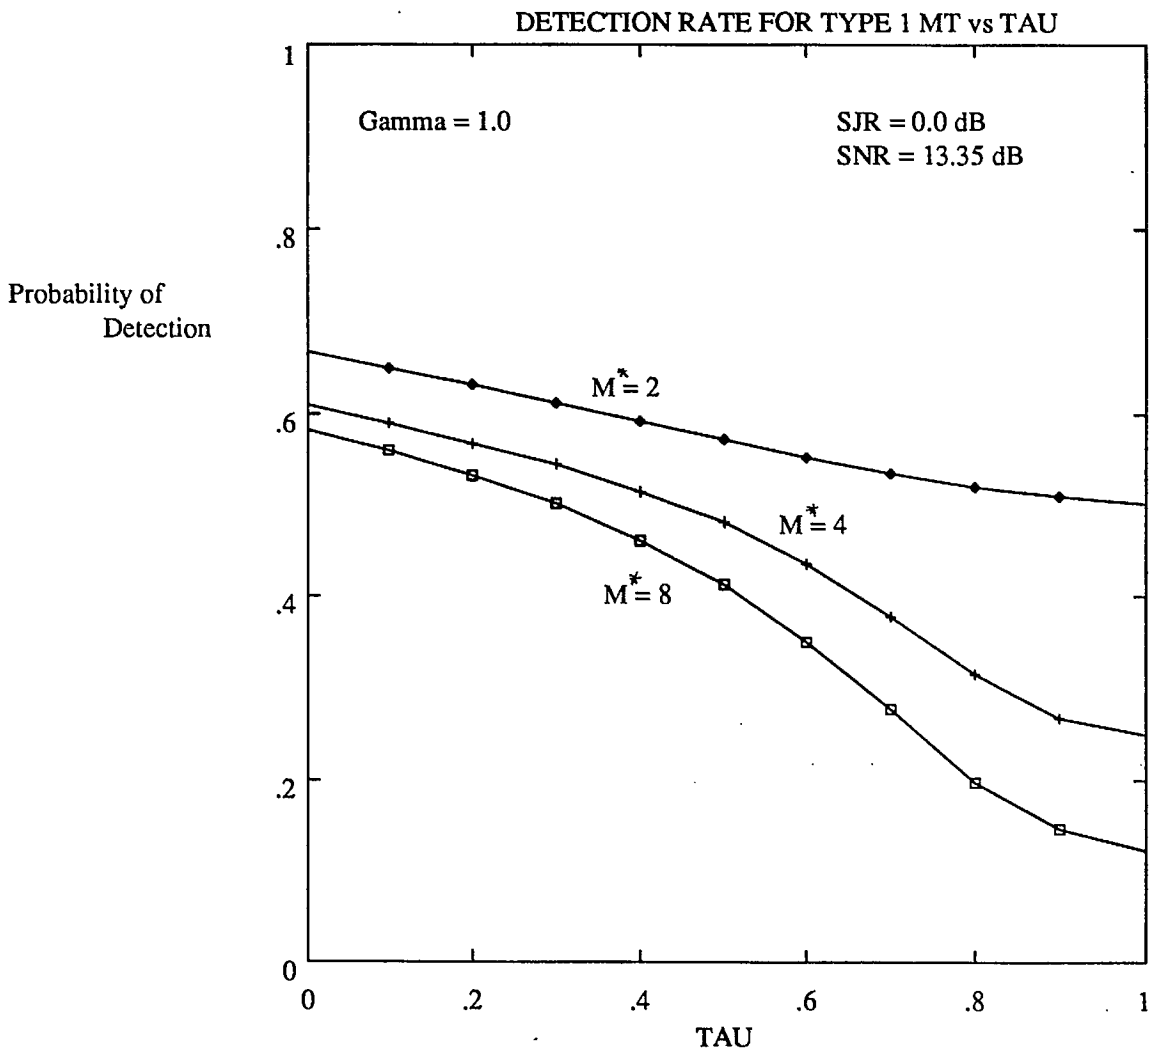

Figure 27: Probability of sync tone detection, $M^* = 2, 4, 8$, SJR= 0.0 dB, $\gamma = 1$.

$M^* = 8$ also has a high absolute P_d, we would certainly expect that using $M^* = 8$ would lead to faster sync with much less chance of false acquistion. However, Figure 28 shows that the jammer can minimize P_d by choosing a lower γ for $M^* = 4, 8$. It is at these lower γ that false acquisitons are more probable. Figure 29 shows probability of sync tone detection for $M^* = 2, 4, 8$ at the γ that gives the worst case P_d for each (found from Figure 28). Now the slope for the $M^* = 2$ curve is maximum, implying that $M^* = 2$ might minimize the probability of false acquisition. However, we know from previous results that $\overline{T_{acq}}$ is too big for $M^* = 2$. The $M^* = 4$ slope here is worse than $M^* = 8$, implying that the $M^* = 4$ case is the harder to handle. We have chosen $M^* = 4$ for evaluating the sync strategy in the next section, as it should give worst-case results for P_{facq}.

We can now comment on the type III MT-jamming case. It should be clear from the previous discussion that the false sync problem is related to the addition of the user sync tone and a jamming tone, and the resulting comparison of total energy to that in the other $M^* - 1$ bins (which may also contain a jamming tone). It is the presence of a jamming tone in one or more of these other bins that causes the problem. For a type-III jammer, however, there can be only one tone bin jammed. If that jammed bin is the sync tone bin, then we are almost certain to correctly detect on that bin. If the jammed bin is one of the other $M^* - 1$ bins, then what happens depends on the relative strength of the jammer and user. If the jammer tone energy is even slightly bigger than the user tone energy (after accounting for the energy loss due to misalignment τ), then the jammed bin is virtually certain to win (given rather modest system background noise levels assumed). On the other hand, if it is even slightly smaller, then we are virtually certain to correctly detect the user tone. In other words, there will be a sharp dividing line between these two eventualities, depending on the exact SJR. This implies that if the jammer levels are low enough that we can sync at all, it will be easy to discriminate between a cell of higher τ from a cell of lower τ, because of the different signal energy levels at these different τ. For this reason, type III MT-jamming should not create a problem with false syncs on partially aligned cells. Either we sync on the best ("correct") cell, or we don't sync at all (if the SJR is too low). For this reason, type I MT-jamming is more problematic, and we have chosen to concentrate on it. Because of the highly predictable nature of the result for type III MT-jamming, we do not explore type III

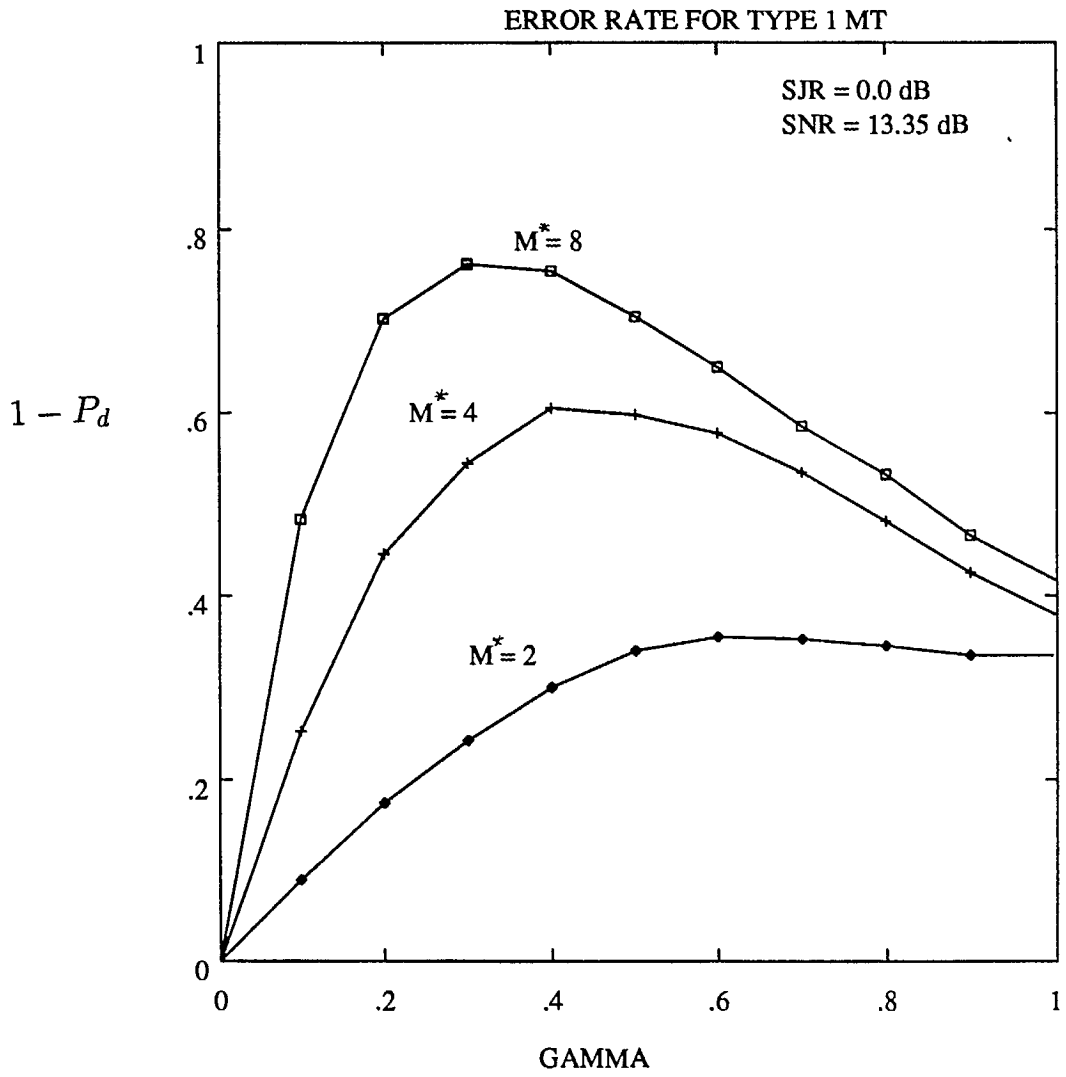

Figure 28: Probability of missed tone detection, $M^* = 2, 4, 8$, SJR= 0.0 dB.

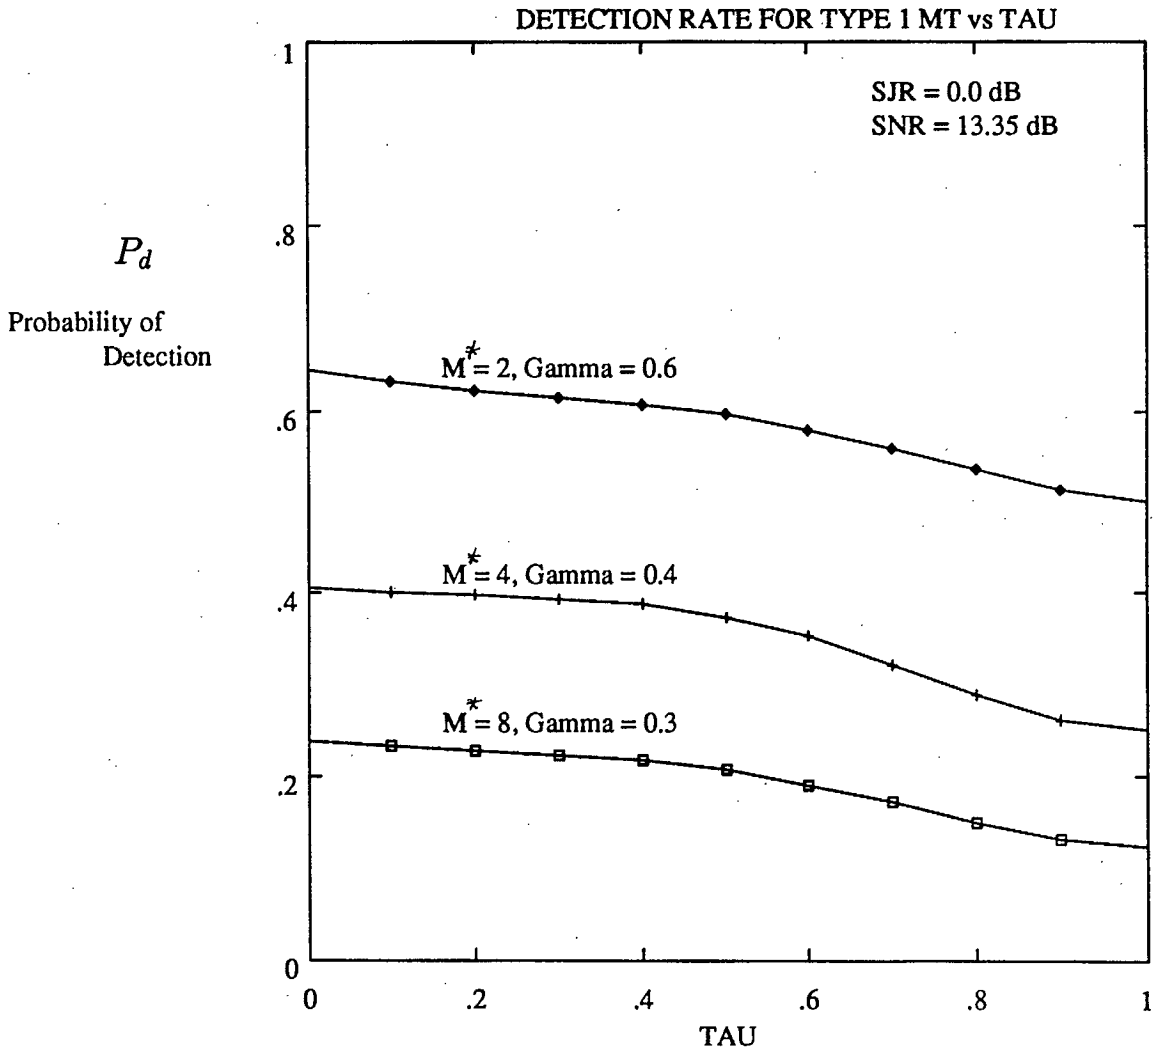

Figure 29: Probability of sync tone detection at worst γ, $M^* = 2, 4, 8$, SJR= 0.0 dB.

MT-jamming in detail in following sections, and instead concentrate on type I MT-jamming.

6.3 Sync Strategy (modification 1)

We modify the strategy presented so far to deal with the false sync problem at SJR = 0.0 dB identified in the previous section. Because of the low differential P_d at the different τ's, we must increase the observation time to make a reliable decision. Now, the steps of the modified procedure are as follows:

1. Make a first pass through all cells with $\Delta T = T_h$, using N_1 hops in each cell, and find the two best adjacent cells (having highest detection counts). The best of these must exceed the count threshold Γ. If it does not, repeat the first pass until it does.

2. Starting from the first of these two cells identified by step 1), construct six cells as below, using steps of $\Delta T = T_h/4$. This ensures that the alignment will be maximized in at least one of the cells.

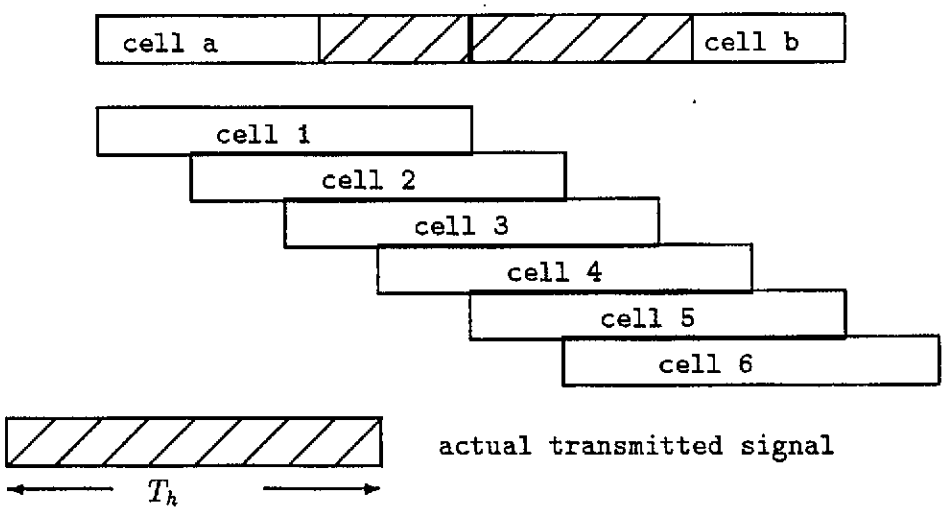

3. Perform a second pass over all these six cells, using N_2 hops of observation in each cell, and choose the cell with the highest number of sync bin detection counts. Declare sync in this cell.

Figures 30- 32 show the resulting residual timing errors at sync for this two-pass procedure. We use type I MT-jamming at SJR=0.0 dB, and the worst-case γ of 0.40. Note that false sync errors ($\tau > .50$) appear if N_2 is not large enough (.e.g. $N_2 = 200$ in Figure 30). With higher N_2, the distribution tightens up towards smaller average residual τ. Because of the limited number of simulation runs (1000), we would not see false sync detections occurring with a probability less than about 10^{-3}. We have not evaluated P_{facq} explicitly, as we have one more modification to make. The problem with this sync search strategy is that the large N_2 required to avoid false sync will lengthen the time to acquisition unecessarily if the SJR is better than SJR=0.0 dB. The solution to this is to make the observation time adaptable.

6.4 Sync Strategy (modification 2)

The modified sync strategy is called a progressive two-pass scheme. It is intended to keep sync times low when SJR is higher than expected. The steps are as follows:

1. Make a first pass through all cells with $\Delta T = T_h$, using N_1 hops in each cell, and find the best cell (having highest sync bin detection count). This count must exceed the count threshold Γ. If it does not, repeat the first pass until it does.

2. Starting from the cell identified by step 1), construct five cells as below, using steps of $\Delta T = T_h/4$. This ensures that the alignment will be maximized in at least one of the cells.

3. Set count-modifier $\rho = 1.0$, set total observation hops $N_T = 0$, set running-total detection counts $C_i = 0$, $i = 1, 5$ for the five cells.

4. Perform a new pass over all these five cells, using N_2 hops of observation in each cell, and add the detection counts to the running totals C_i. Find the cell with highest C_i

5. Set $N_T = N_T + N_2$, set $\rho = \rho - 0.10$

6. If the best cell has a detection count in excess of ρN_T, declare sync and exit, otherwise, go to step 4).

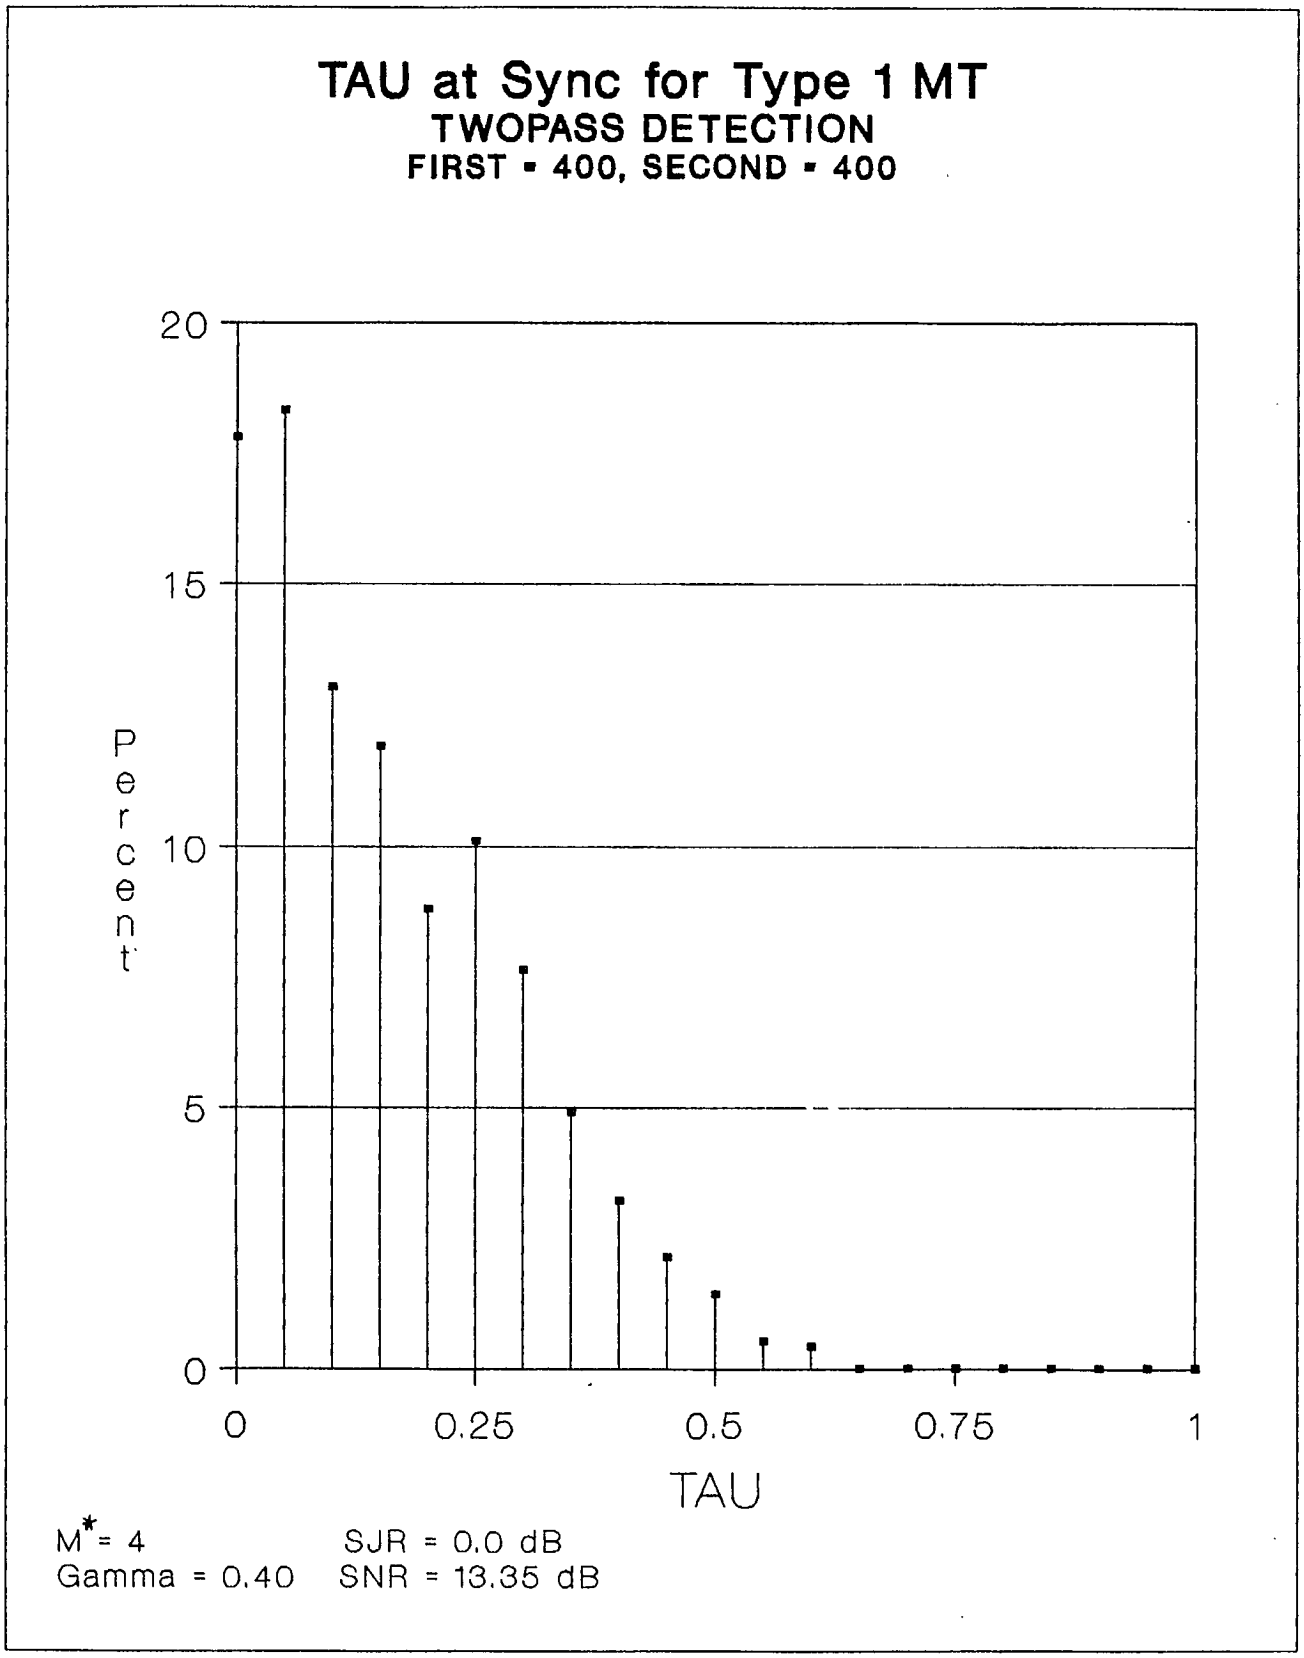

Figure 30: Residual timing error distribution at sync.

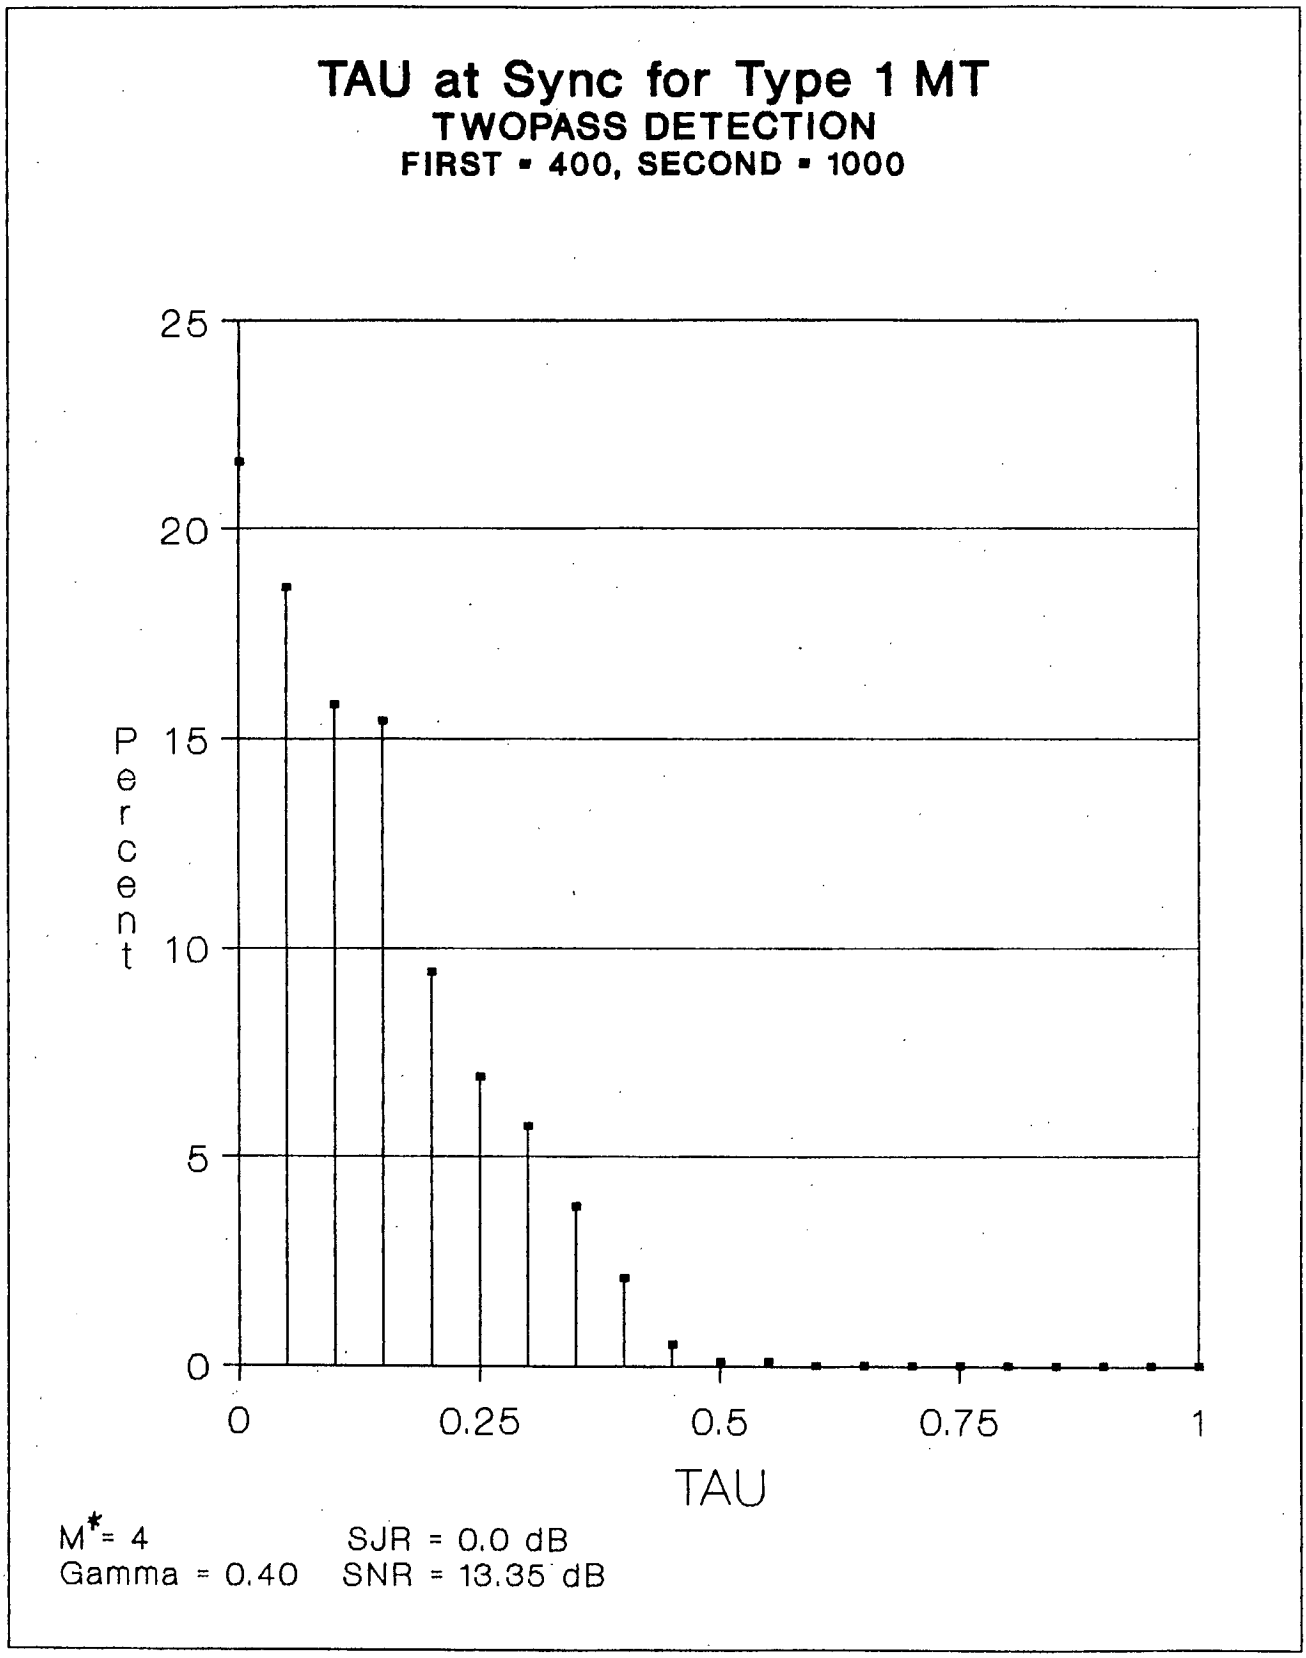

Figure 31: Residual timing error distribution at sync.

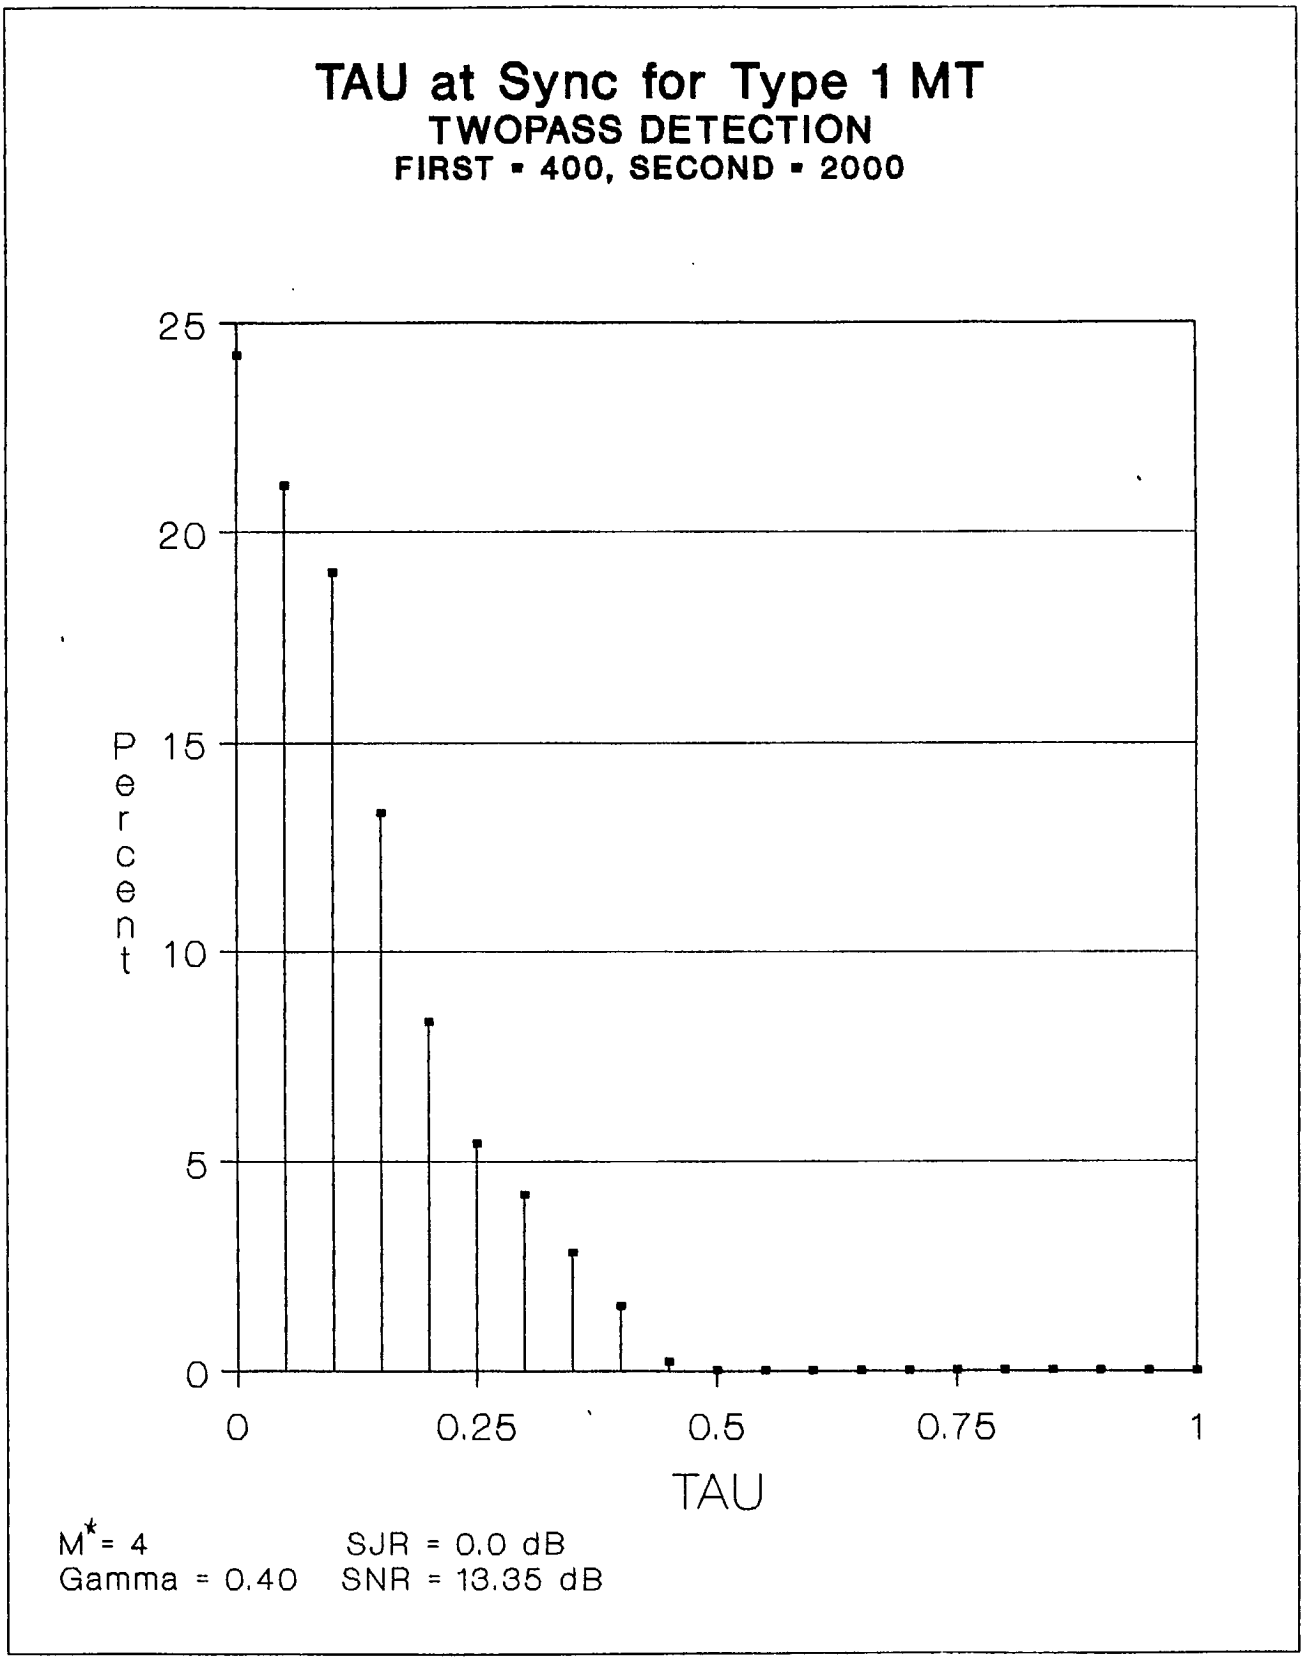

Figure 32: Residual timing error distribution at sync.

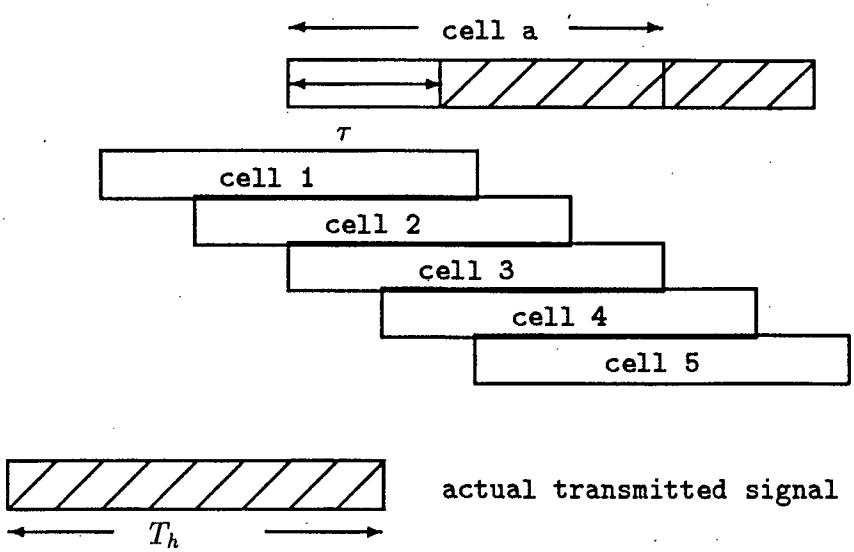

This procedure accumulates observation time as successive passes are made, and lowers the acceptance criterion (by reducing ρ) on successive passes to ensure convergence. At high SJR, sync will be declared quickly after the first pass over the five cells. At lower SJR, up to 10 passes will be performed (at which point $\rho = 0$, and we simply pick the best cell). We keep N_2 not overly large in order to reduce sync time in the event of early sync. Note that this scheme can be fully pipelined (as in Figure 7) for the transmissions of the second and subsequent passes.

Figures 33- 42 show distribution of residual timing errors at sync, and corresponding distribution of time-to-sync, for this new procedure. Again we consider type I MT-jamming at SJR=0.0 dB, over the range of $\gamma = 0.20$ to $\gamma = 1.0$. The first pass uses $N_1 = 400$ with a threshold $\Gamma = 150$, and subsequent passes over the five cells constructed after the first pass use $N_2 = 200$. Mean acquisition times are all on the order of 1 second. Although γ around 0.60 causes the highest $\overline{T_{acq}}$, the residual timing errors are greatest around $\gamma = 0.4$; this is also where the probability of sync error is greatest (residual $\tau > 0.5$) as expected from earlier Figure 21. Although we do not have exact results on this P_{facq}, worst-case probabilites are presented in Table 2. By worst-case, we mean the alignment case that causes the best cell to have $\tau = 0 + \epsilon$, and one of the other four cells to have $\tau = 0.50 + \epsilon$ (for $\epsilon \rightarrow 0$). This particular alignment minimizes the difference in the detection probabilities P_d for these

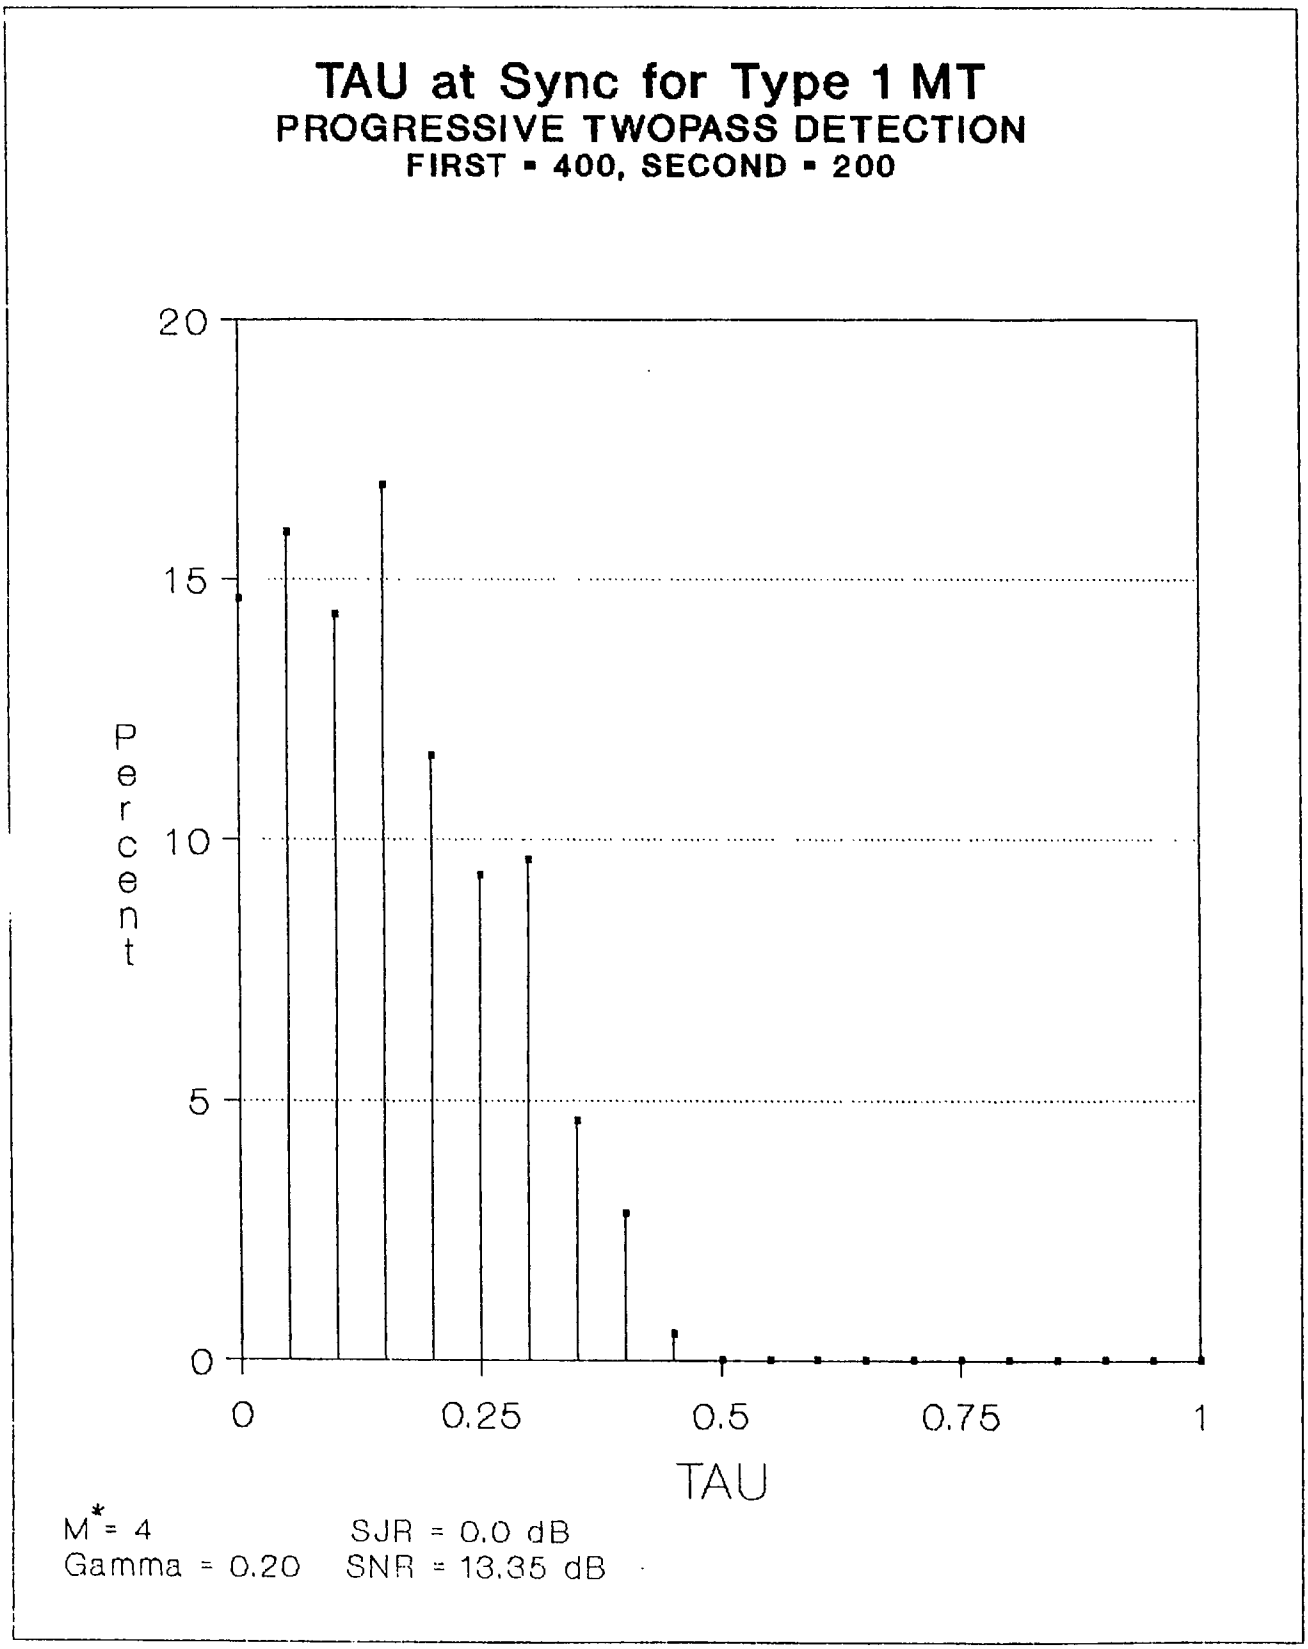

Figure 33: Residual timing error distribution at sync.

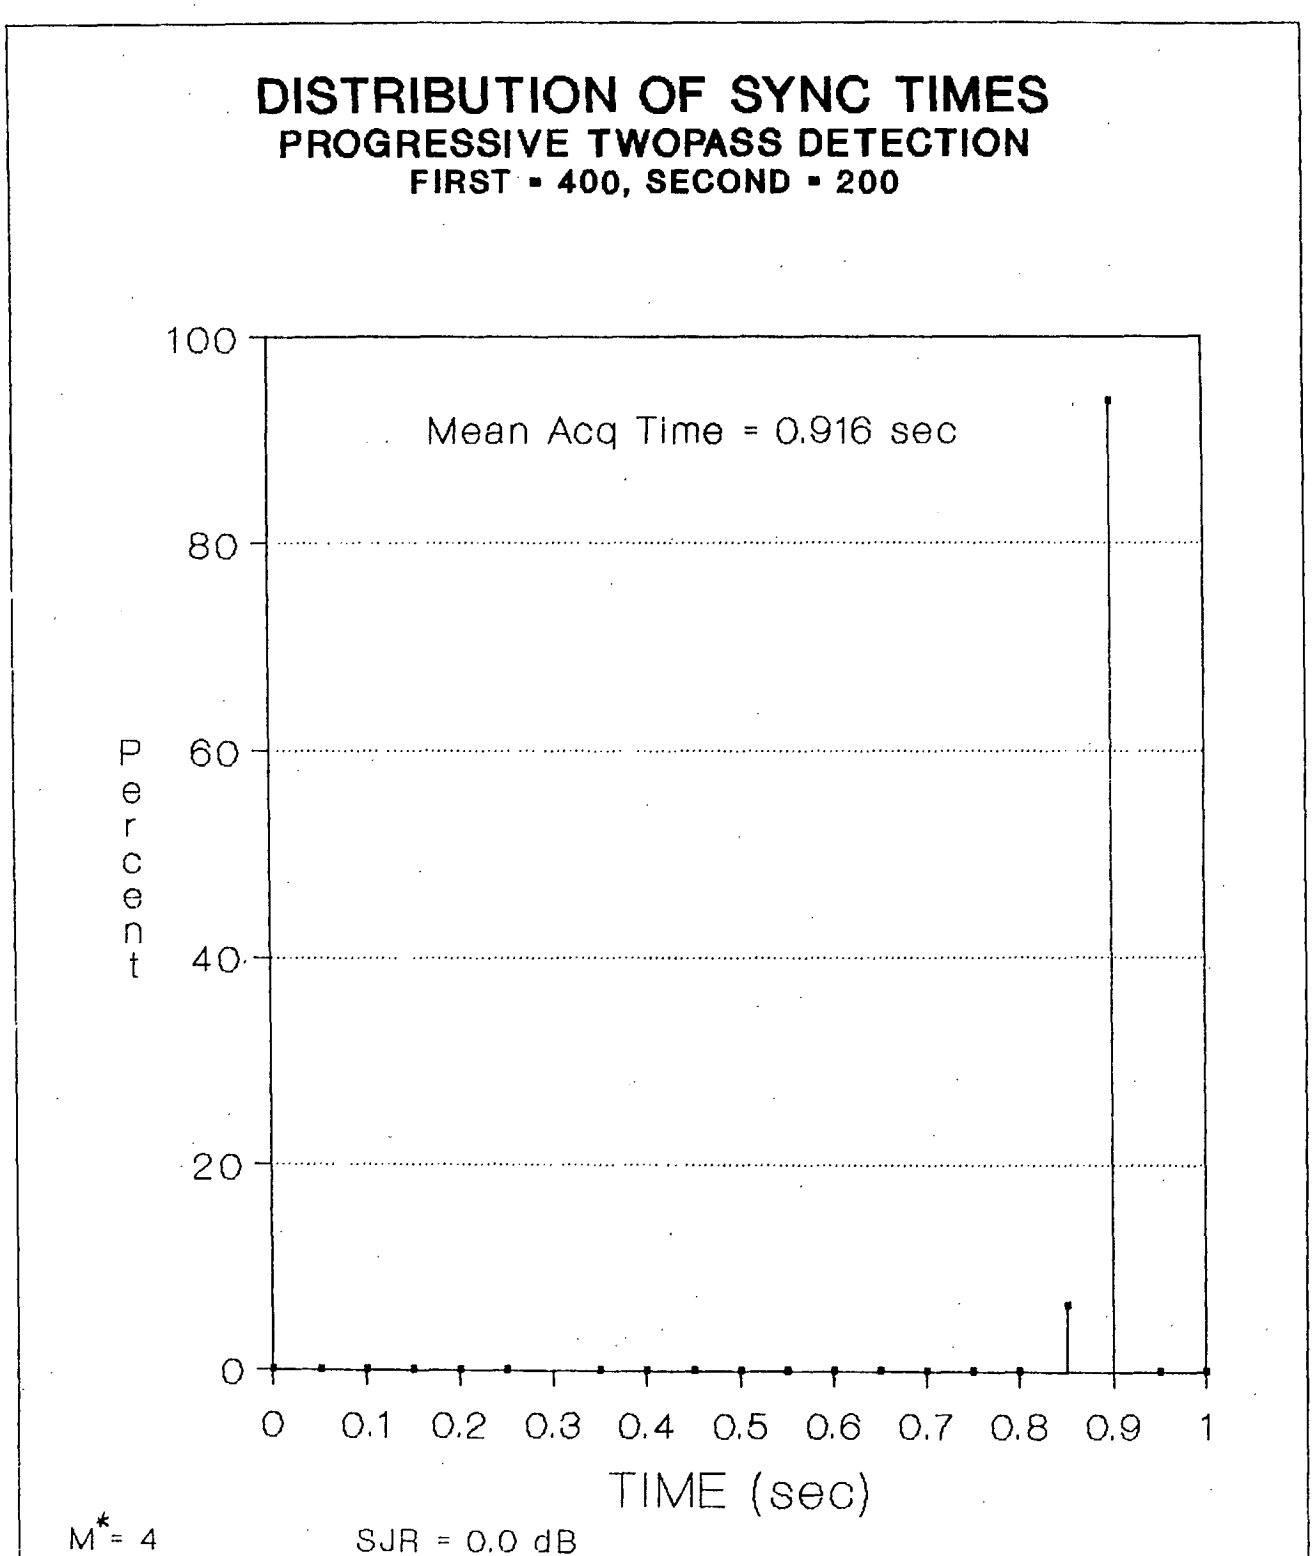

Figure 34: Times to sync.

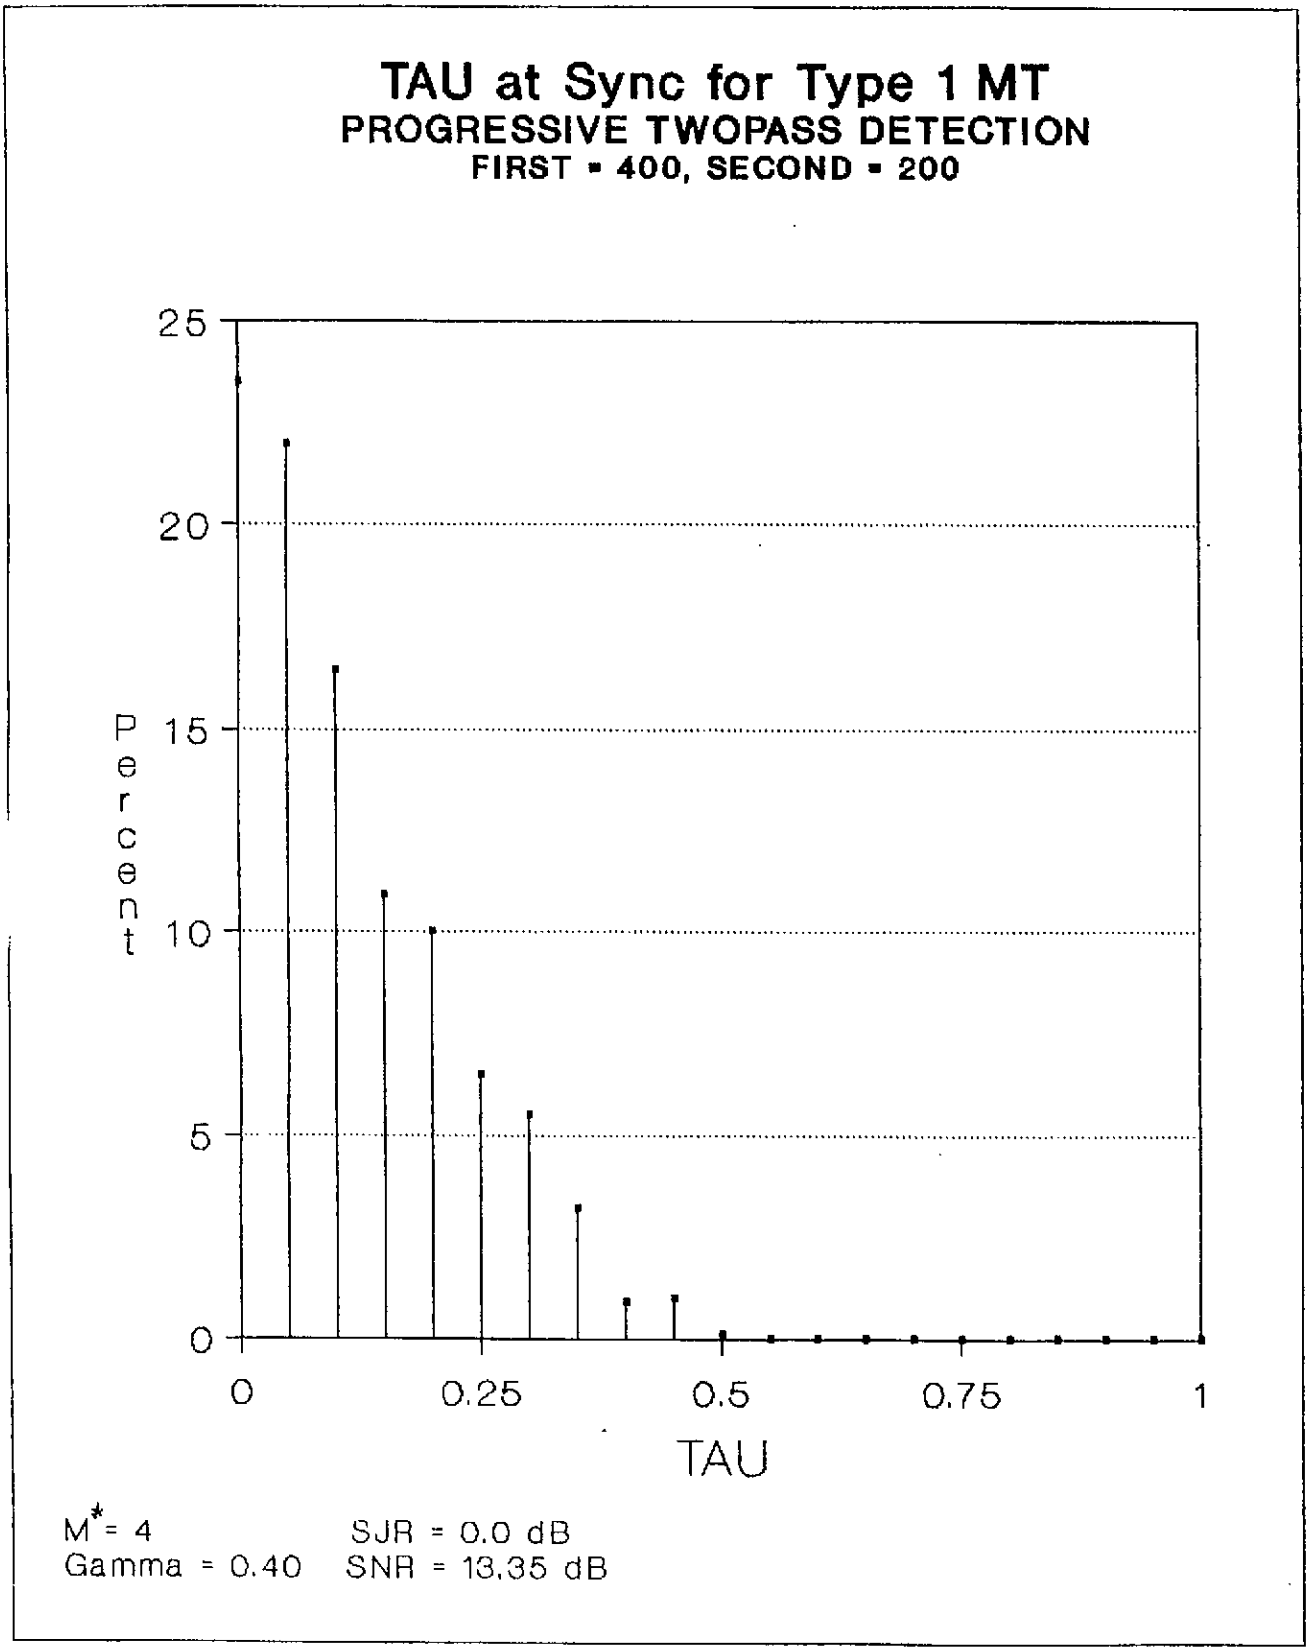

Figure 35: Residual timing error distribution at sync.

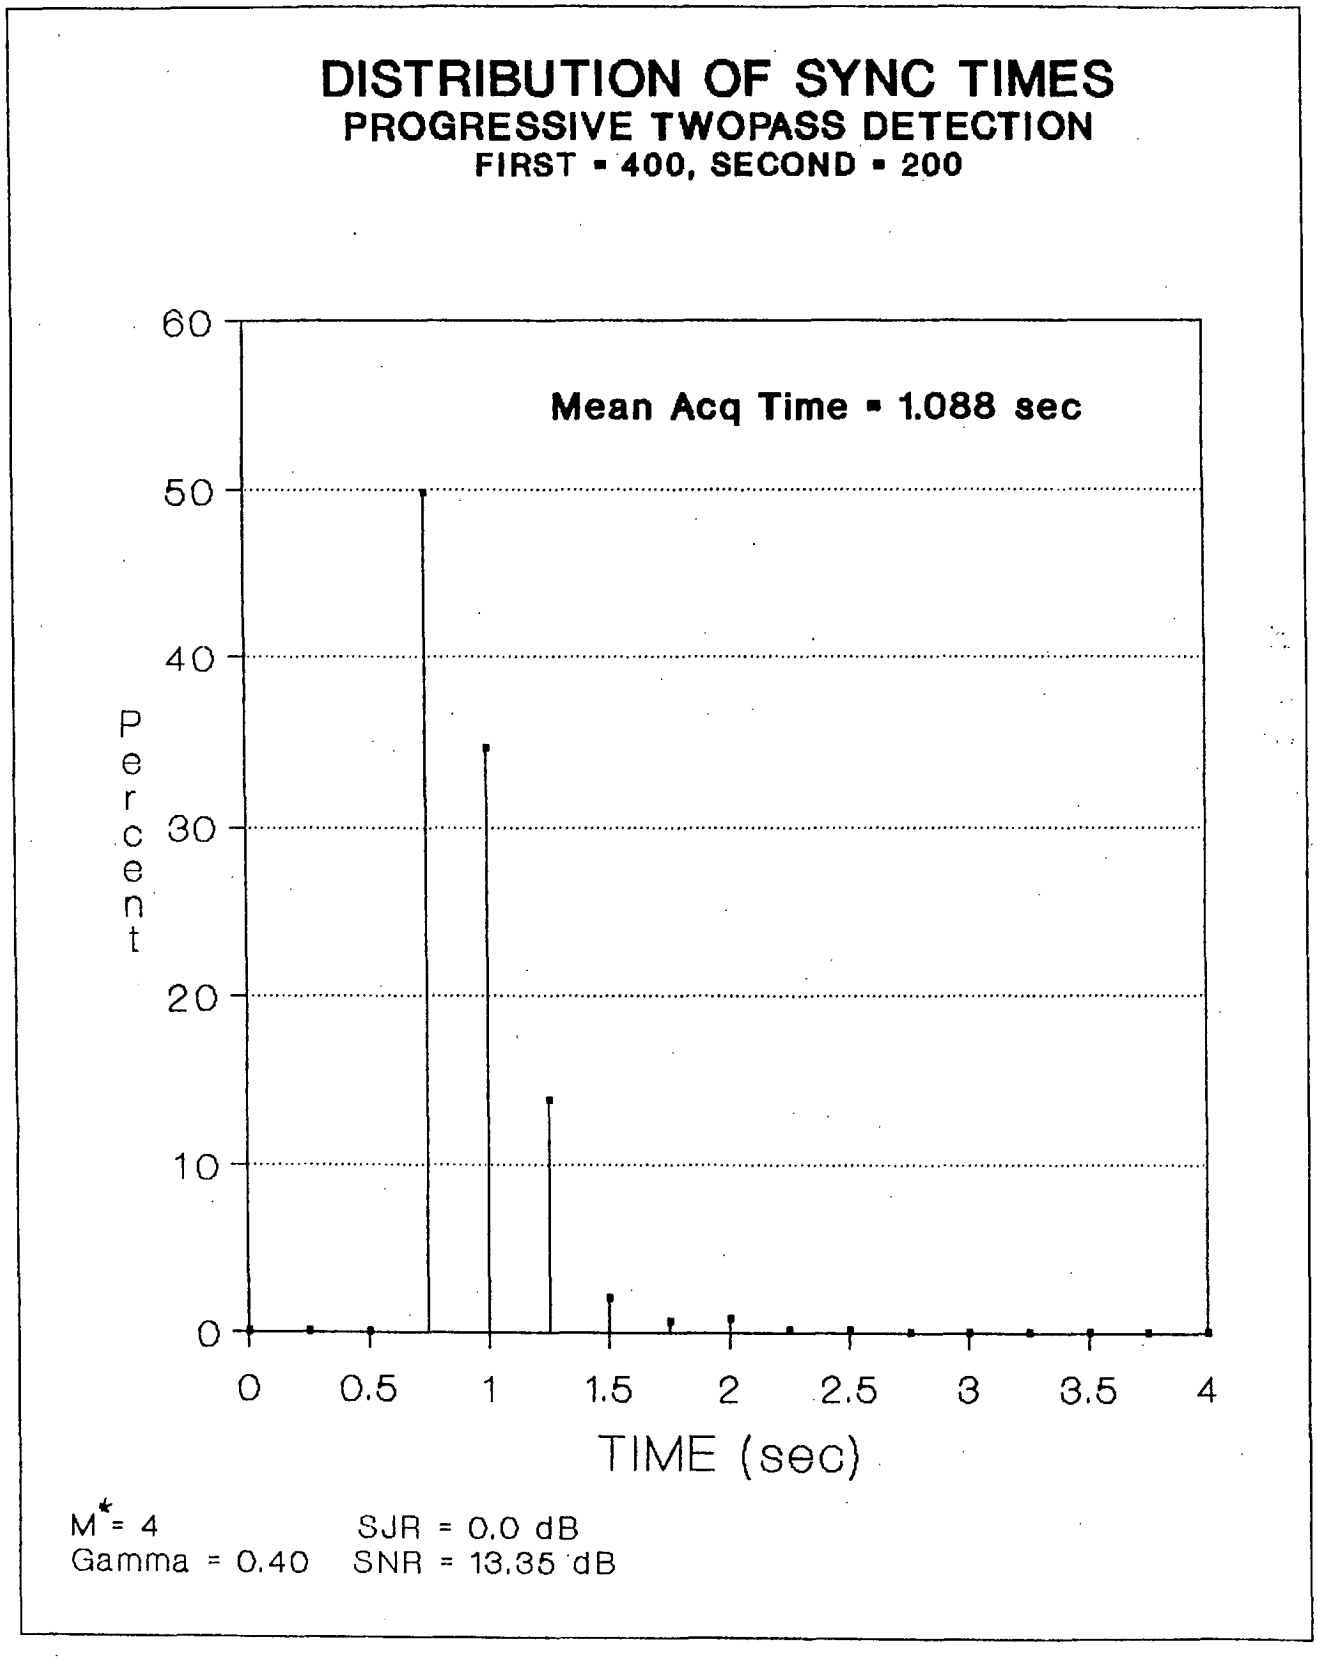

Figure 36: Times to sync.

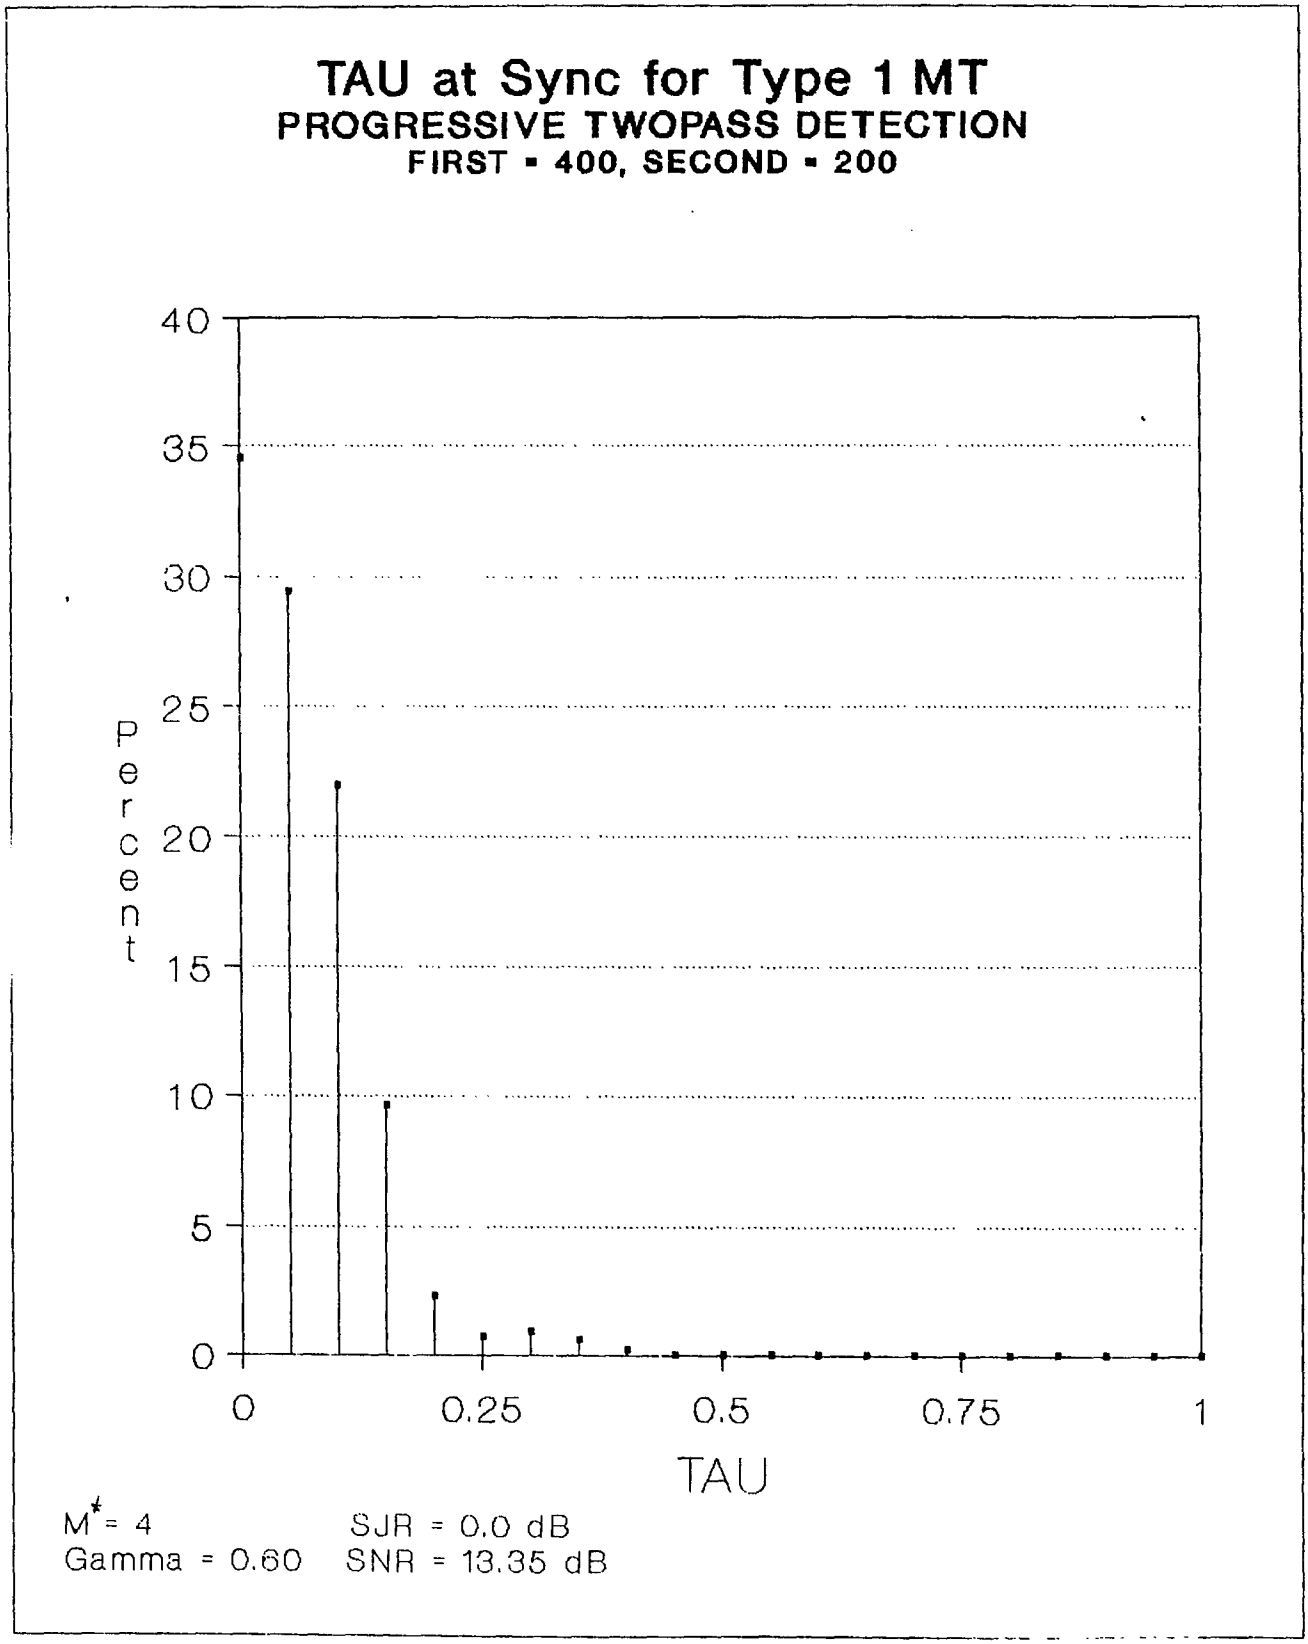

Figure 37: Residual timing error distribution at sync.

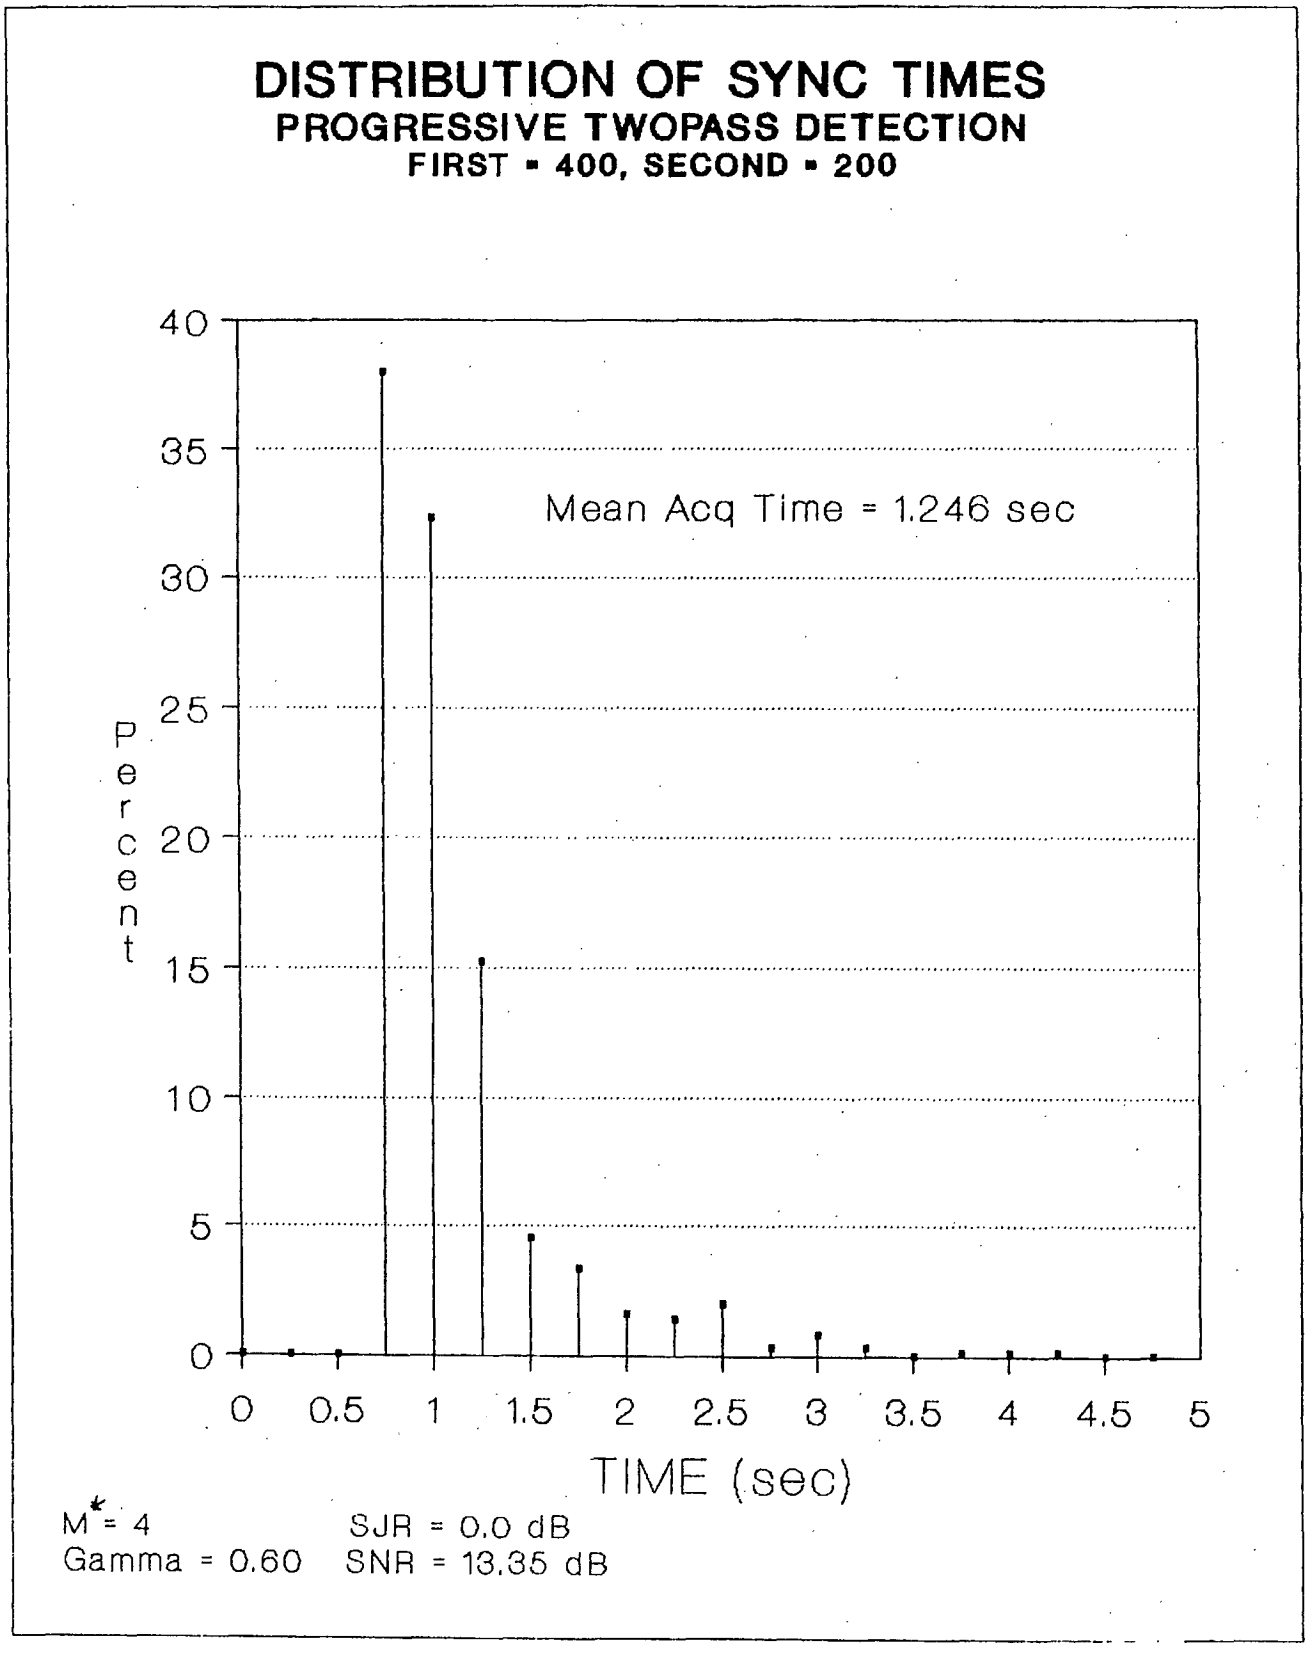

Figure 38: Times to sync.

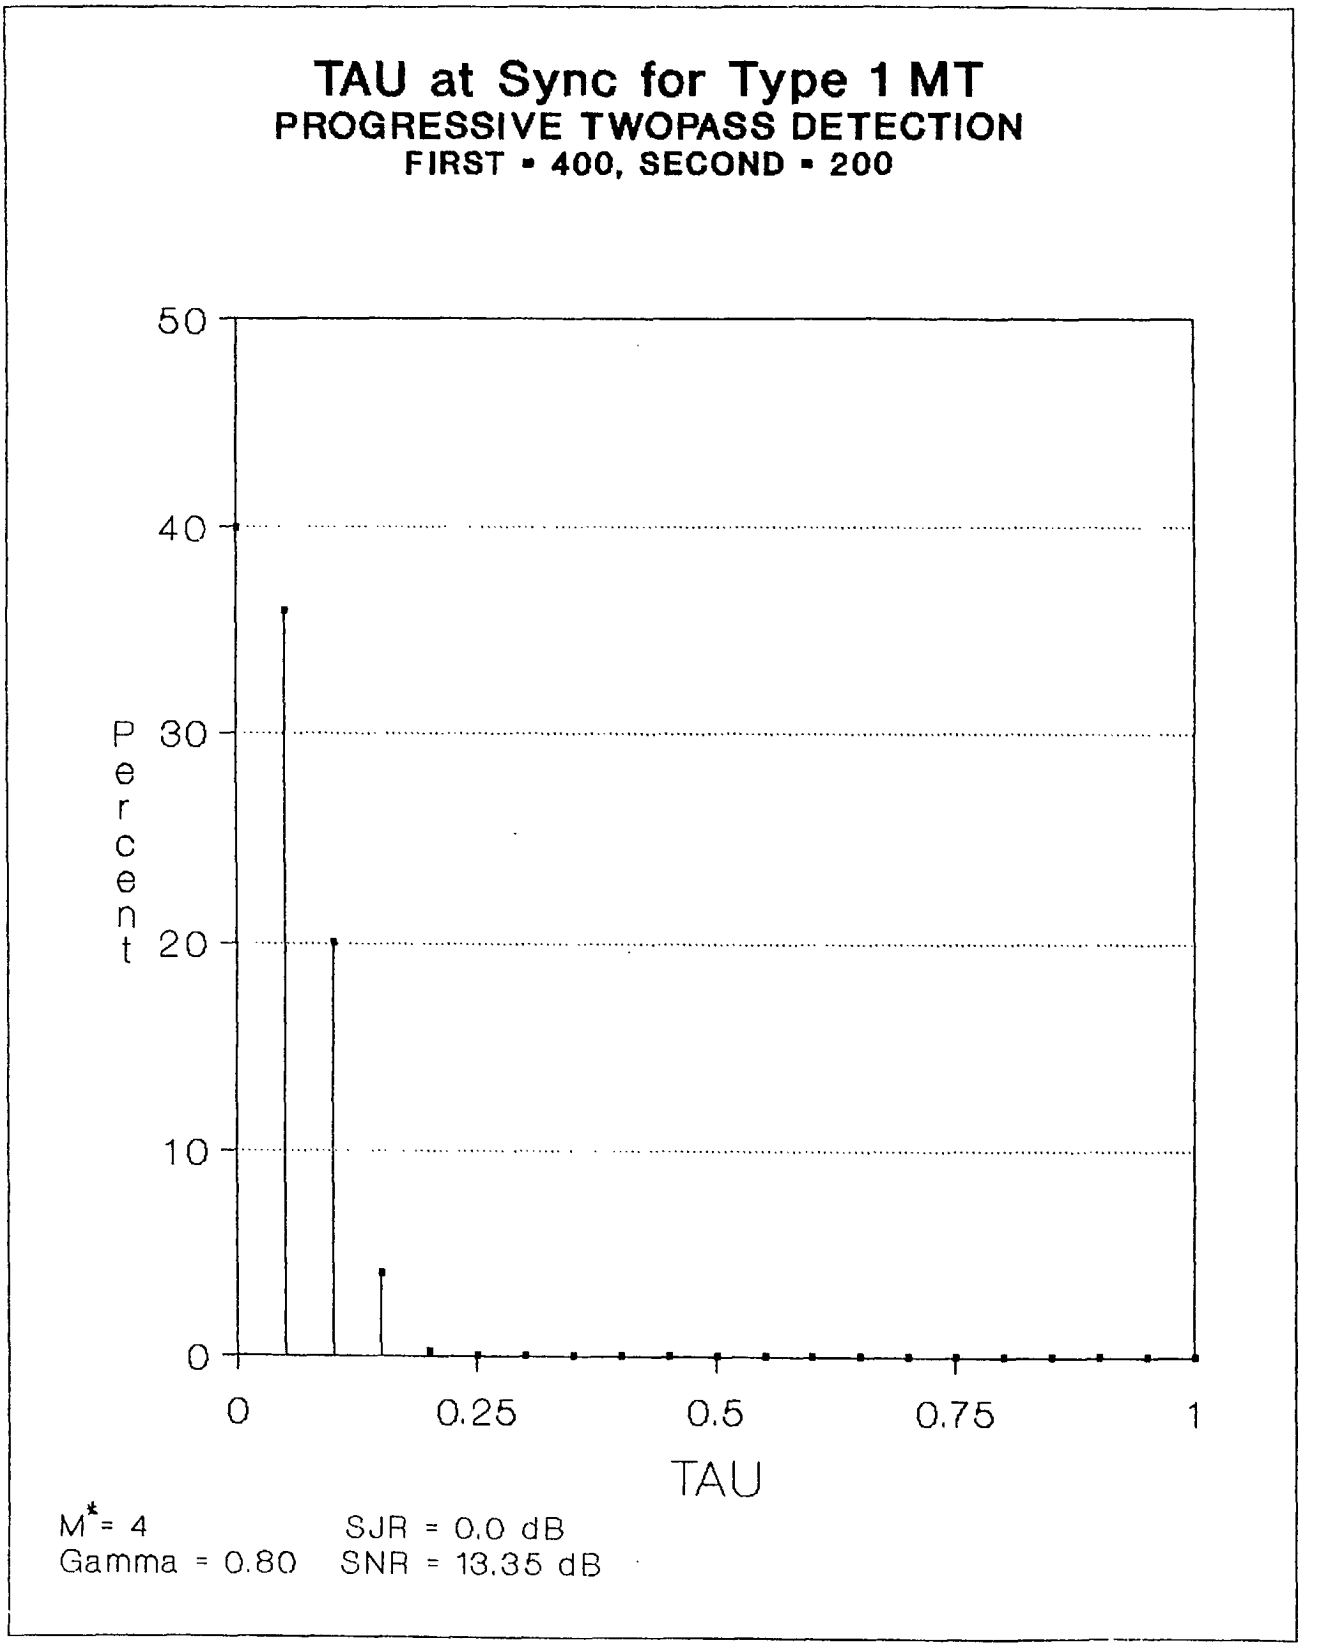

Figure 39: Residual timing error distribution at sync.

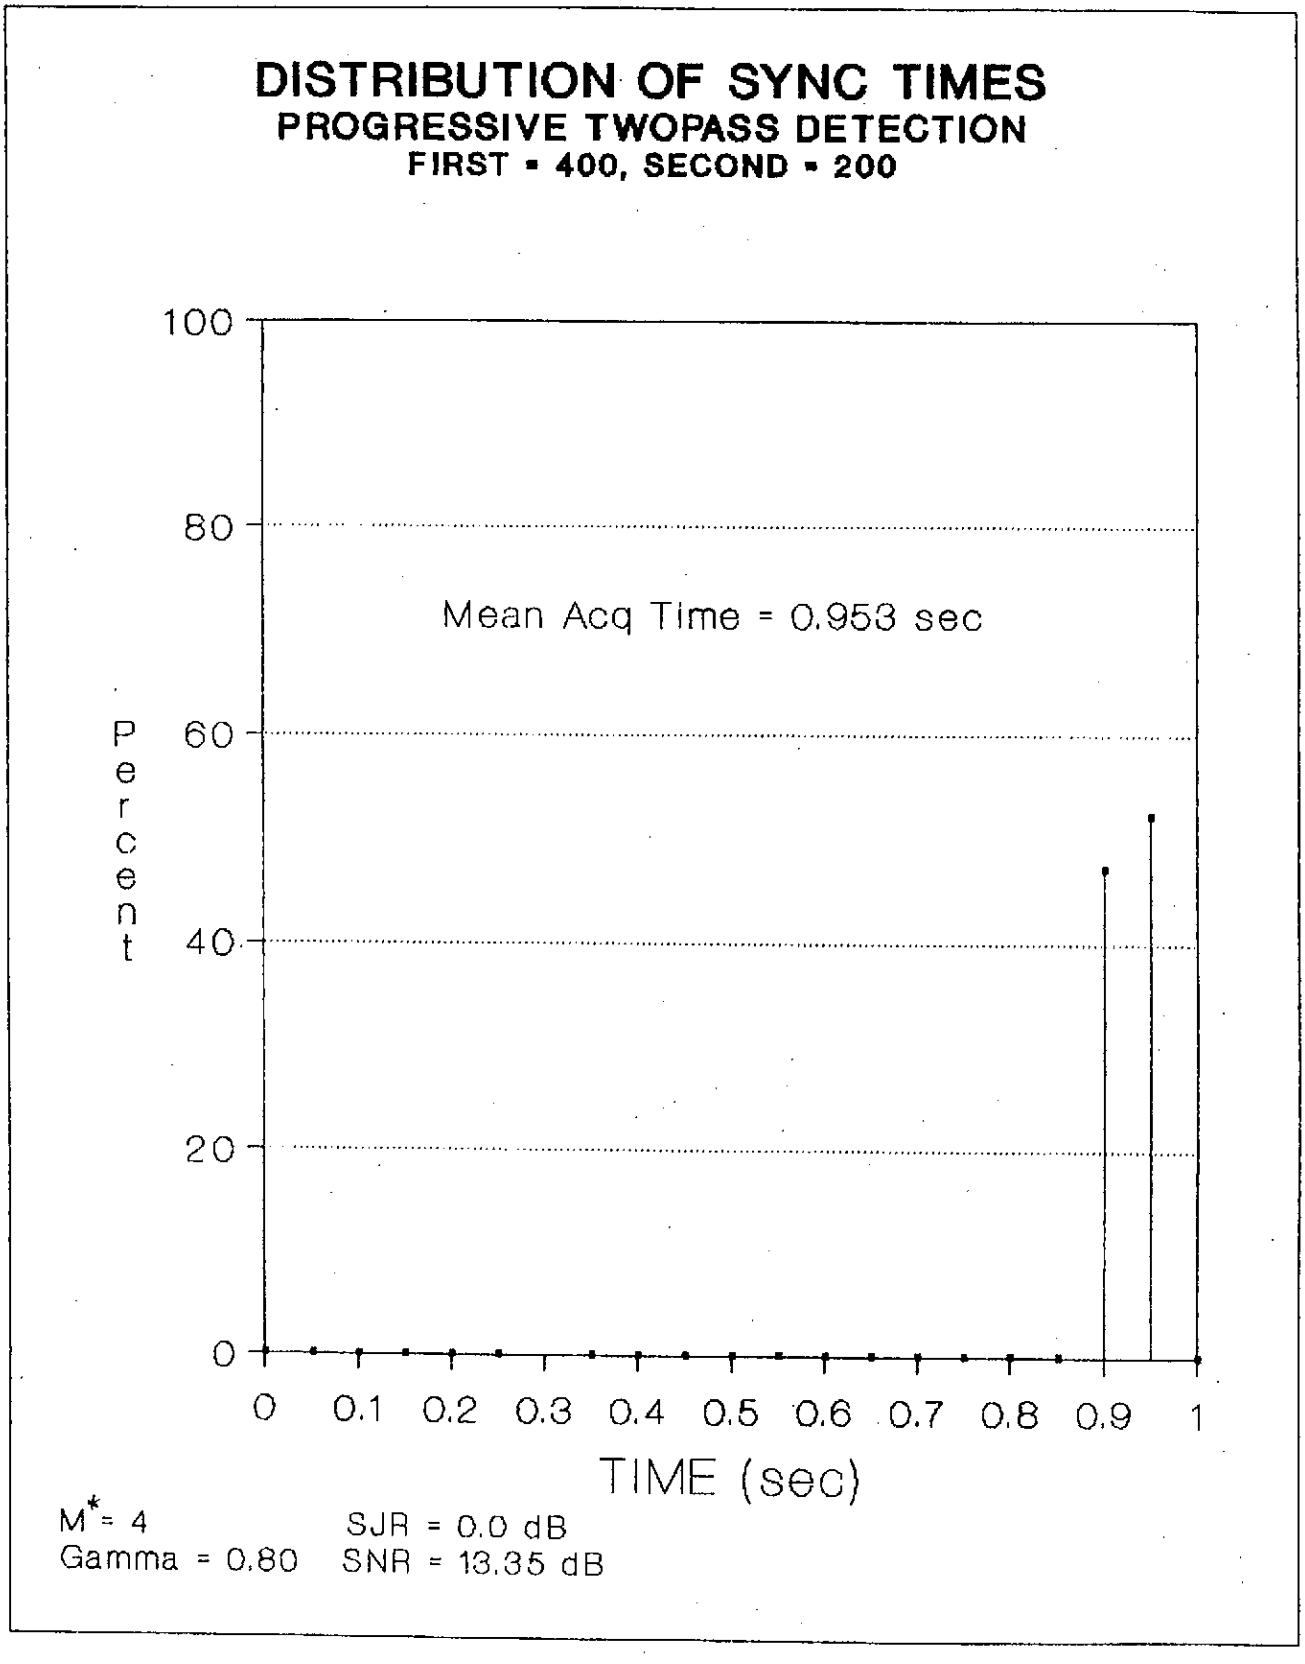

Figure 40: Times to sync.

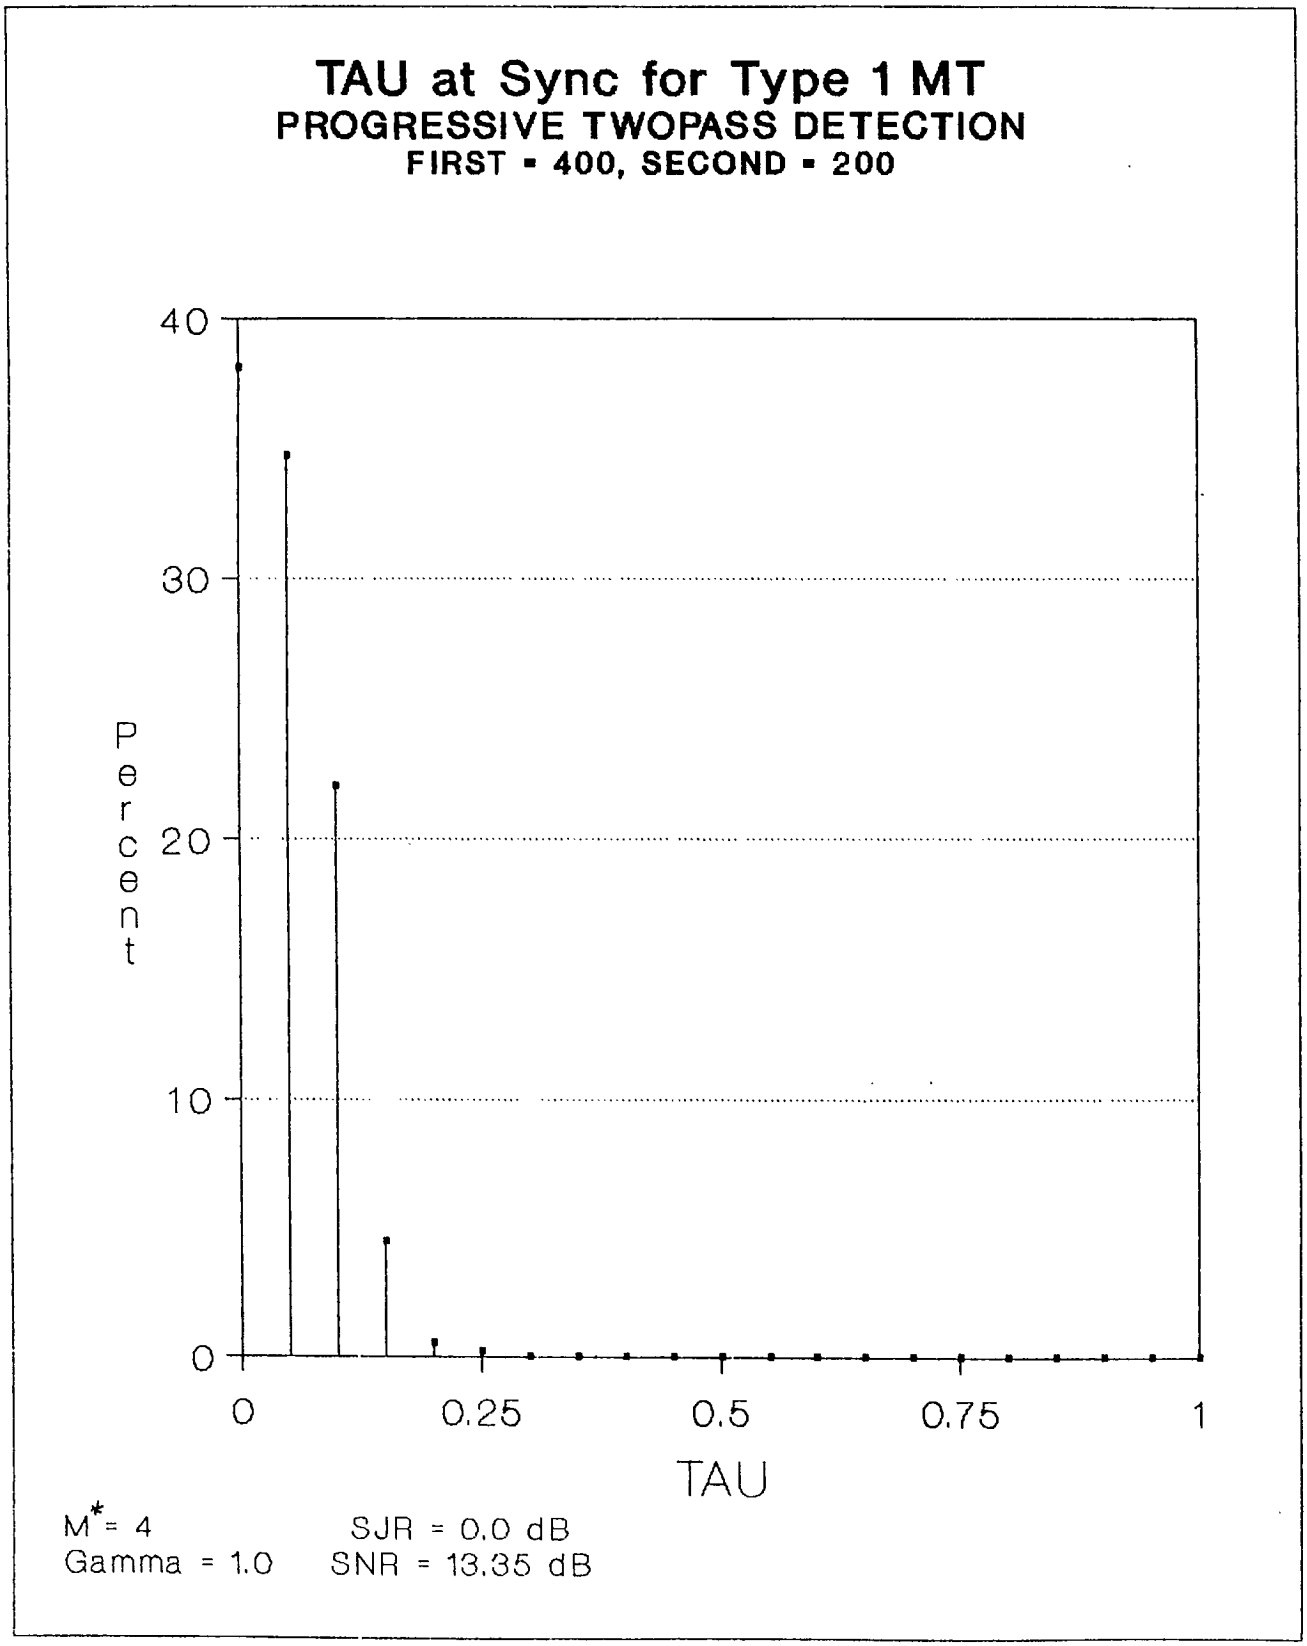

Figure 41: Residual timing error distribution at sync.

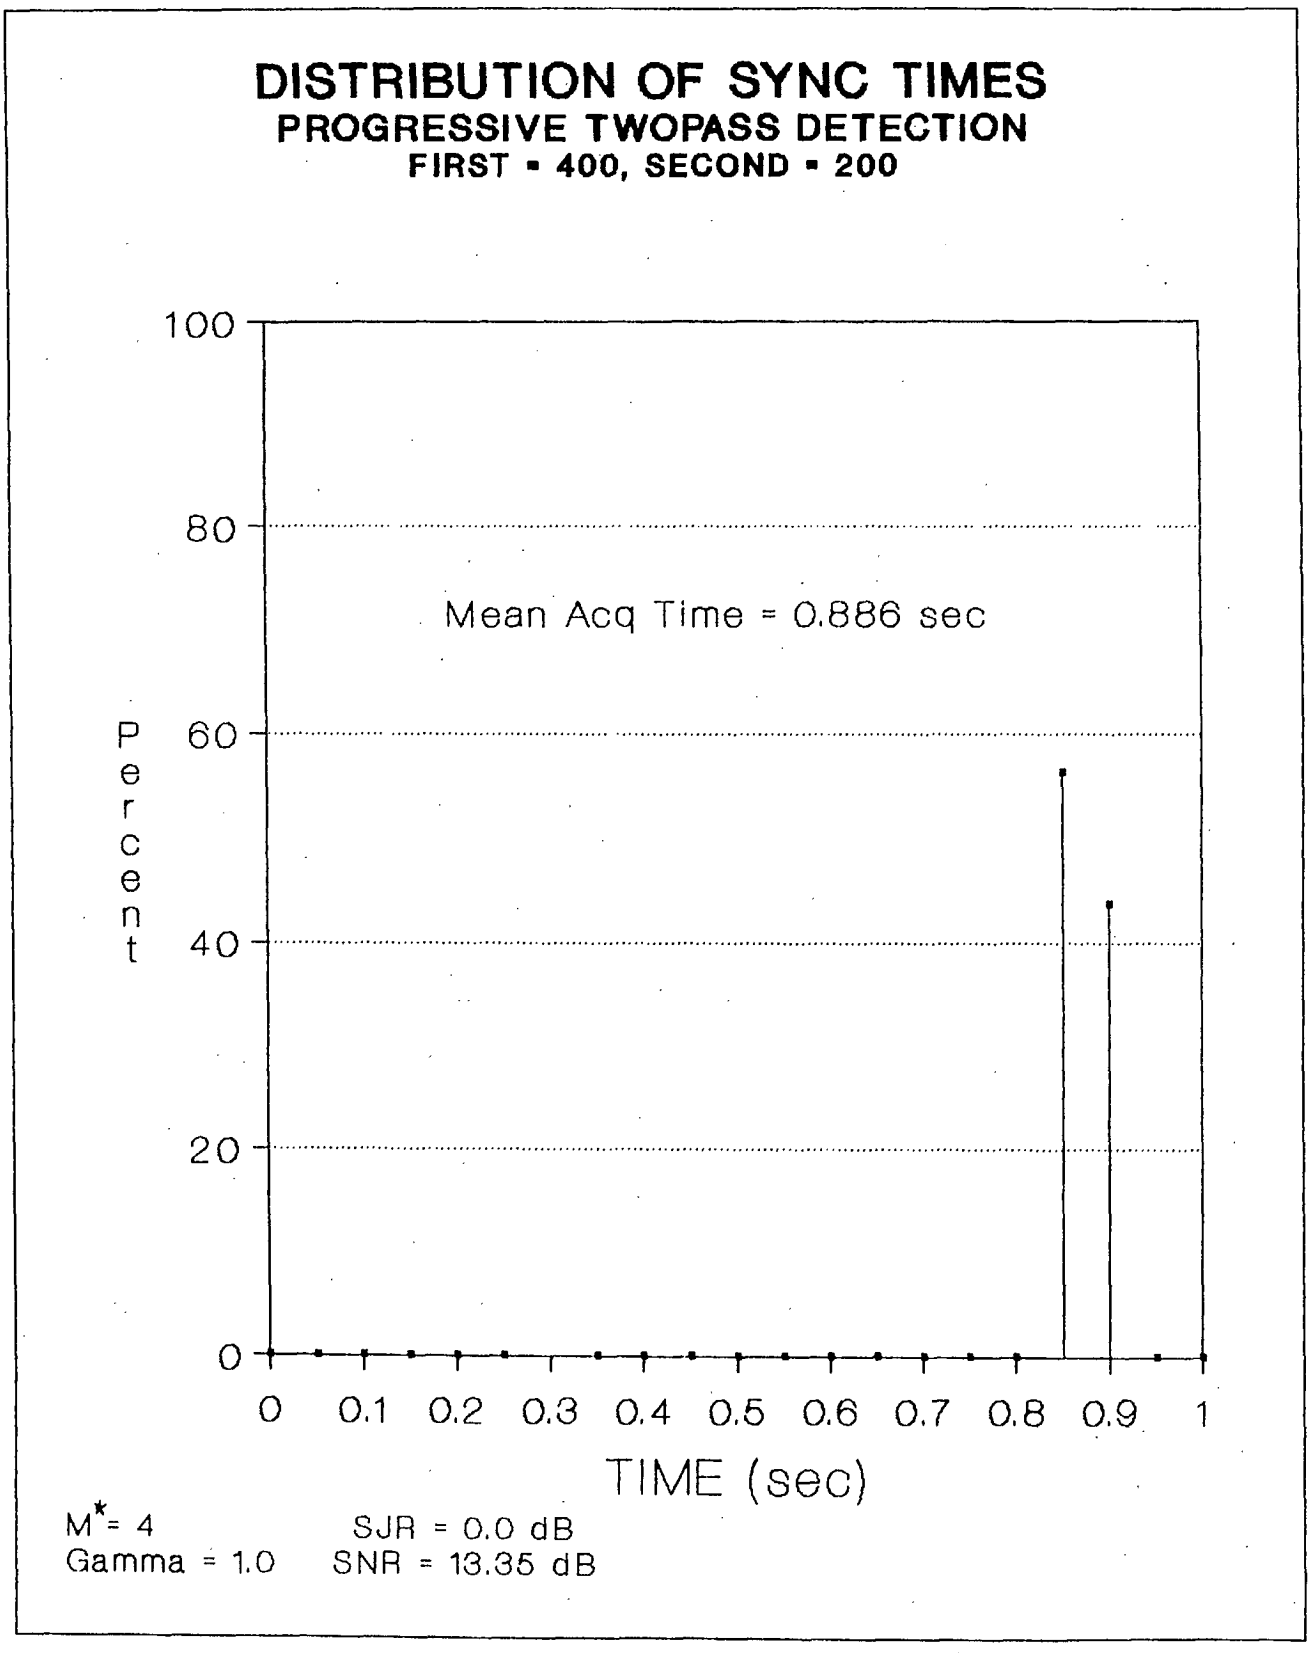

Figure 42: Times to sync.

γ	$\overline{T_{acq}}$	P_{facq}
0.20	0.916	2.9×10^{-2}
0.40	1.088	1.0×10^{-2}
0.60	1.246	1.0×10^{-6}
0.80	0.953	1.4×10^{-13}
1.00	0.886	2.2×10^{-18}

Table 2: Upper bounds on P_{facq} for worst-case alignment.

two cells (see Figure 24), maximizing the probability of incorrectly declaring sync on the incorrect ($\tau = 0.50 + \epsilon$) cell. Of course, once we average over all possible random alignments, the net P_{facq} will be smaller.

For higher SJR's, (Figures 43- 44 , $\overline{T_{acq}} = 0.86$ and 0.68 seconds, respectively), acquiring sync is relatively easy, resulting in low residual timing errors, and a decline in P_{facq} to levels less than 10^{-6}.

We have verified expected behaviour with type III MT-jamming. Our results indicate that type-III jamming allows much easier discrimination between cells with different τ, so that P_{facq} is easily kept below 10^{-6}, if SJR is sufficient to allow sync at all.

6.5 Sync Strategy (modification 3)

There is one fairly trivial problem with the sync procedure as defined so far. In the event of low, or very low jamming levels, any of the five cells defined after pass 1 that contain even a portion of the signal energy will easily pass the subsequent threshold tests, giving no criterion for choosing between these cells. The same effect occurs for type III MT-jamming with a low fraction of the band jammed β, regardless of SJR. The procedure in the previous section actually depends on jamming and background noise to help differentiate between cells having different energies due to differing alignments τ. It stops at the first cell that passes the threshold criterion, and, in the event of low jamming and noise, may inadvertently

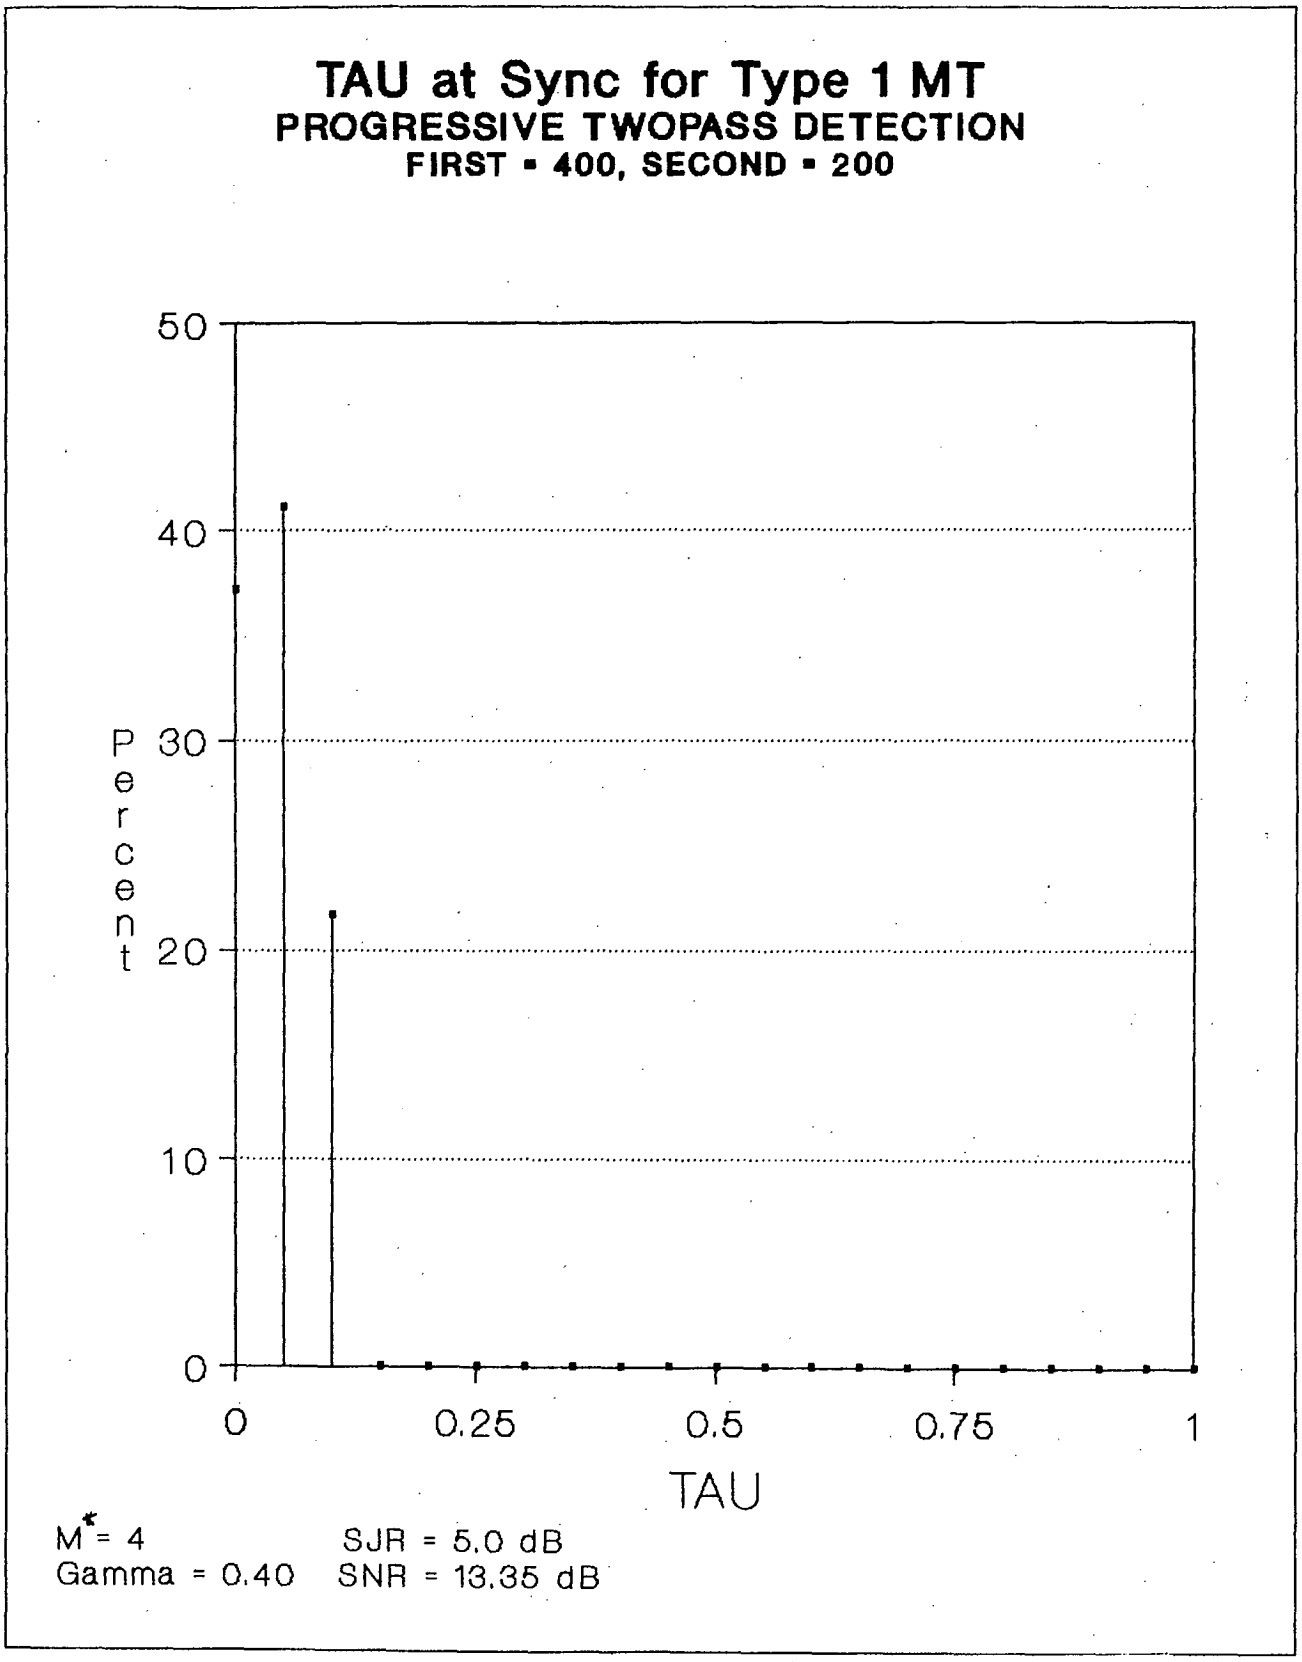

Figure 43: Residual timing errors at sync, SJR = 5.0 dB.

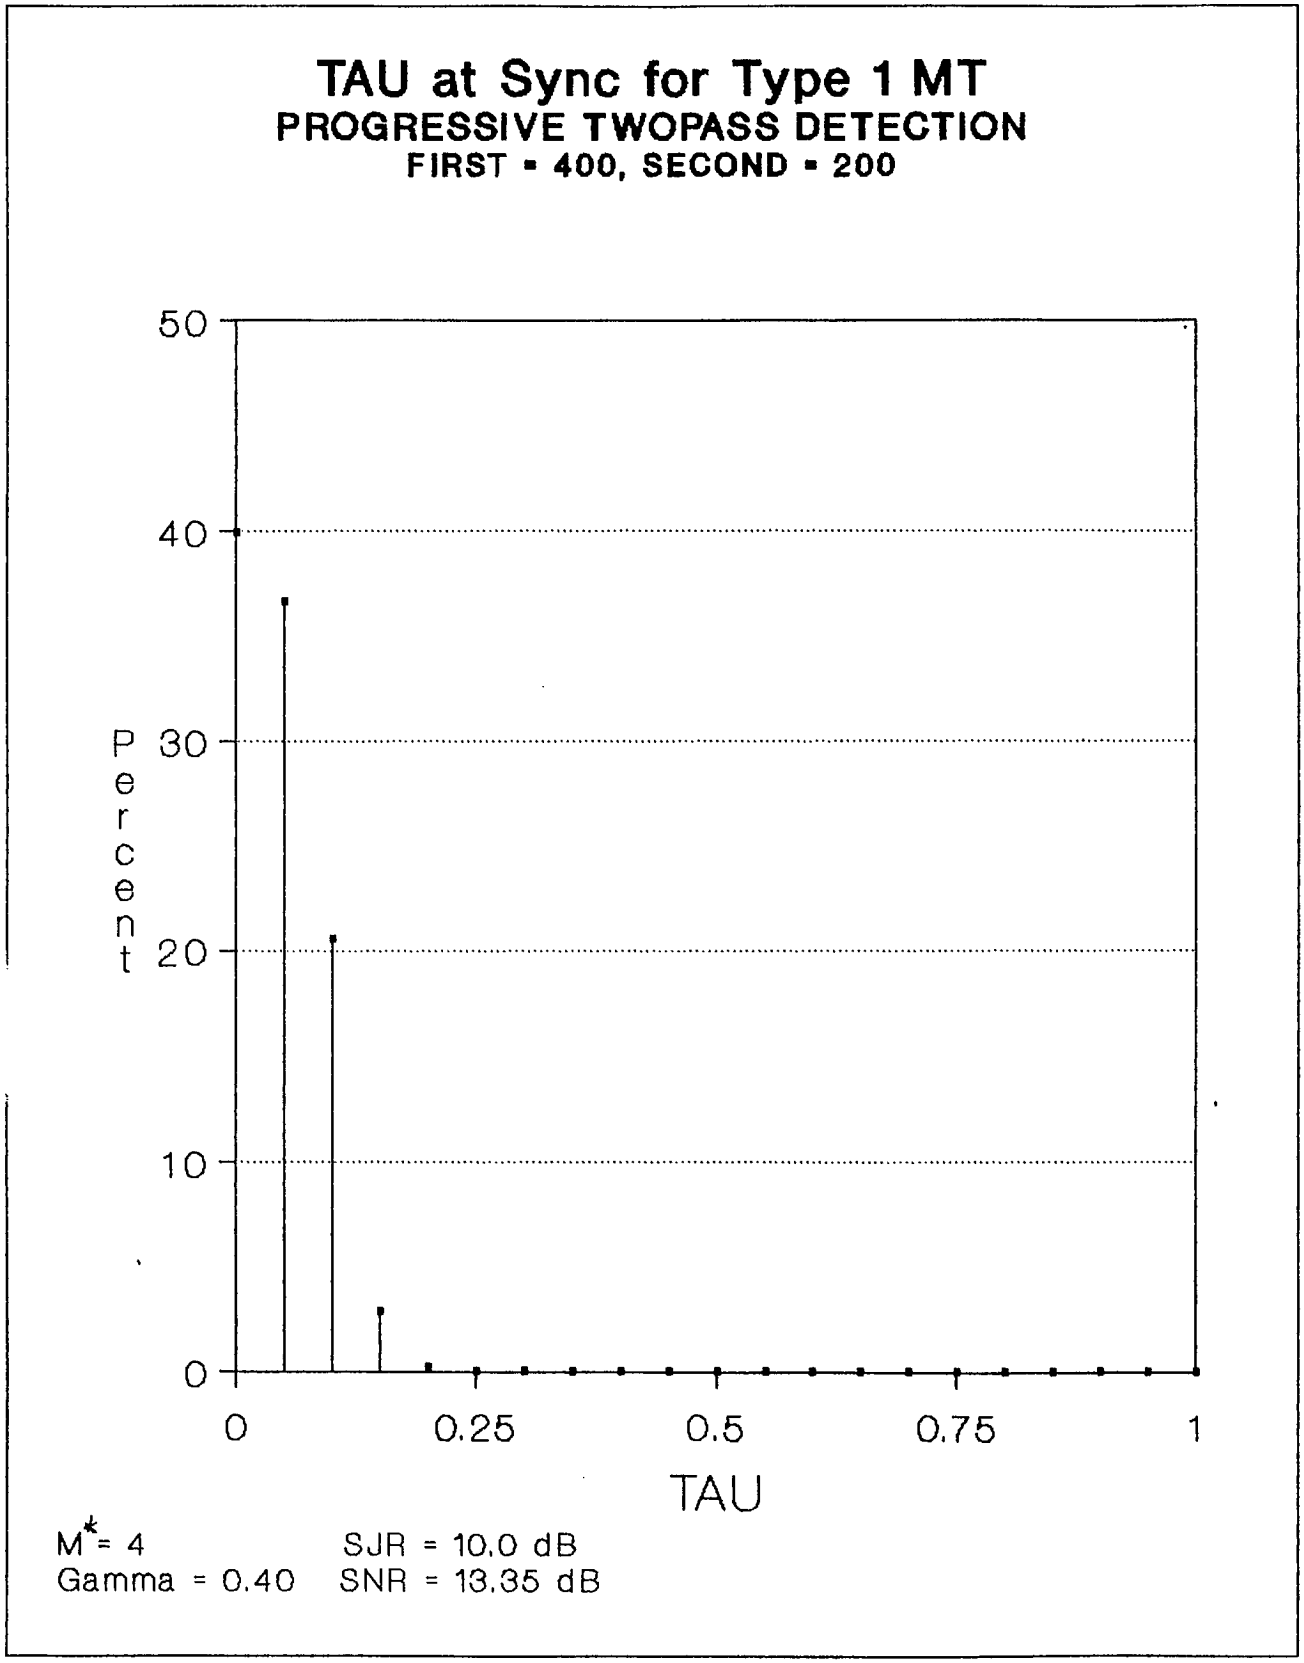

Figure 44: Residual timing errors at sync, SJR = 10.0 dB.

declare sync on a cell having a high τ (although not one having $\tau = 0$). One solution is to evaluate all five cells at the current threshold, and then declare sync in the cell having the highest detection count. This should work well at all but exceedingly low jamming or noise levels. In the event of detection-count ties (or near ties) in a contiguous group of cells in the second pass, we can simply declare sync at an alignment half-way between the first and last cell in this group. It can be seen (most easily for the case of zero jamming and zero noise) that this last procedure is guaranteed to declare sync at an alignment that has $\tau < 0.50$. Of course this also means that our residual τ's may be actually *worse* when jamming and noise are very small. The only way around this is to use some measure of the actual energy levels that are detected in the sync tone bin at the satellite, in addition to the hard decsions assumed for our sync procedure. Actually, the probability of ties given the size assumed for N_2, and given a background noise SNR of 13.35 dB, is so small that such anomalous behaviour is not expected. We should therefore not need detected-energy information to be provided by the satellite .

Note that we had not included this modification for the results presented in the previous section. However, the jamming levels considered there were more than high enough that declaration of sync in a cell of $\tau > 0.5$ was being caused by missing the best cell due to its high jamming levels, rather than by the fortuitous absence of jamming making more a poor cell (that happened to be checked before the best cell) pass the threshold. In other words, we do not expect the results presented to change significantly with the incorporation of these modifications.

7 Summary

We have proposed and investigated a simple strategy for acquiring coarse synchronization. We have shown that sync can be acquired quickly (on the order of 1-2 seconds for the initial ± 5 hops of initial timing uncertainty assumed) even at very high jamming levels. In the absence of sync errors, residual timing errors will be less than $T_h/2$, and average around $T_h/4$. Probability of false sync detection (.i.e. timing error in excess of $T_h/2$) is negligible

($< 10^{-6}$) at higher SJR's, and can be expected to remain at levels well below 10^{-3} at the highest jamming levels considered (more detailed calculations are required to determine these levels more exactly). This means that users can expect to successfully synchronize within a second or two at least 999 times out of 1000 attempts.

Multi-tone (MT) jamming with more than one jamming tone possible in a user channel (type I) is found to be more problematic than multi-tone jamming with a single jamming tone per user channel (type III). Although type III MT-jamming has an energy advantage by a factor M over type I, performance is quite predictable (either SJR is sufficient to allow rapid and accurate sync, or it is too low to allow sync at all). Type III can only hamper sync at high fractions of the band jammed. Type I MT-jamming is more effective at lower fractions of the band jammed, reducing its energy disadvantage relative to type III. Performance in the presence of partial band noise jamming is expected to be qualitatively similar to type III MT-jamming, but less detrimental for the same SJR. Synchronization with full band noise jamming of level SNR=0.0 dB is not a problem.

Because of the nature of the satellite-based system and resulting uncertainties in received user signal levels, a traditional synchronization approach based on an energy detection threshold seemed undesirable. Instead, a simple, robust, and transparent approach based on the use of a single sync tone per hop, and best 1-of-M^* tone detections at the satellite, was adopted. The probability of falsely detecing the sync tone on any one hop is independent of jammer power or strategy. In addition, no attempt is made to identify jammed hops. Both of these features lead to robustness. M^* should be at least 4 to minimize acquisition times. Use of diversity combining for detection decisions at the satellite has been discussed, and it was argued that such combining is information lossy for sync purposes, and is best avoided. One of the attractions of the proposed approach is its transparency, that is, it can use the same satellite receiver hardware and downlink resources as needed for normal user transmissions. It shows that we need not incorporate features on the satellite specifically for synchronization purposes. In fact, sync procedures can be further refined *after* satellite hardware design has been completed by modifying the search strategies (for example, in the event that initial uncertainty regions are larger than considered in this study).

The initial region of uncertainity that must be searched by the sync procedure was determined from typical positional uncertainties for a geostationary satellite, given that the internal state of the satellite hopping sequencer is periodically transmitted (in encrypted form) in a common downlink information slot. User transmitter frequency errors were also assumed negligible. The sync search procedures were tailored to deal with the rather small resulting number of cells that had to be searched. Large uncertainity regions may require more complex strategies, such as making a first pass over all cells with small observation time in each, then concentrating on the cells with highest detection counts for further observation (with the option of backing up to repeat the first pass in the event that none of the cells examined in the second and further passes look good). The situation of non-negligible initial errors in user carrier frequencies was also considered. A sync-band approach that promised to greatly reduce acquisition times in the face of large initial carrier frequency uncertainties was presented in Appendix A.

References

[1] P.J. McLane, P.H. Wittke, and S.J. Simmons, "Final Report: A Study of Space Communications Spread Spectrum Systems: Part III Synchronization", D.S.S. Contract no. 36001-6-3530/0/ST, Feb. 1988.

[2] M.K Simon, J.K.Omura, R.A. Scholtz, and B.K. Levitt, "Spread Spectrum Communications: Vol. III", Computer Science Press, 1985.

[3] F.J. Harris, " On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform", Proc. IEEE, Vol. 66, No. 1, Jan. 1978, pp. 51-83.

[4] J.S. Bird and E.B. Felstead, "Antijam Performance of Fast Frequency Hopped M-ary NCFSK - An Overview", IEEE Journal on Selected Areas in Communications, vol. SAC-4, pp. 216- 233, March 1986.

[5] L.J. Mason, "A Method for the Precise Synchronization of a Frequency-Hopped Spread Spectrum System", Technical Memorandum DSAT #7/88, Nov. 1988.

[6] L.J. Mason, "A Second-Moment Method for the Precise Synchronization of a FHSS", Technical Memorandum DSAT #11/88, Dec. 1988.

[7] D.P. Kolba, "Generalized control and networking for EHF Satellite Communication Systems", AIAA 9th Communications Satellite Systems Conference, San Diego, CA., March 1982.

A Appendix A

The presence of large initial frequency offsets in the FDMA group may require a user position "permutation" strategy to be employed to avoid a neighbour's synchronization attempts consistently interfering with an active user (or another user also attempting to synchronize). By permuting user positions within the group from hop to hop, users attempting to synchronize with large initial frequency offsets will appear as random tone jammers to the other active users. The unattractive alternative to this is to restrict synchronizing users to low duty-cycle transmissions.

If large initial frequency offsets are indeed a reality, it is frequency uncertainty that could greatly dominate the time to acquisition of sync. This being the case, it may be worthwhile contemplating dedication of a larger portion of satellite resources for synchronization purposes. An approach that uses a dedicated "sync" band that is shared on a random access basis is presented in this Appendix.

Figure 45 shows the situation with frequency alignment errors. If these alignment errors are less than f_s , we have approximate alignment. All randomly selected sync tones will fall into one of the bins in the user channel that the satellite defines; there will be an equal signal energy loss for each which depends on the residual alignment errors δf and τ defined earlier. If the frequency errors are several times the bin spacing f_s , only a fraction of the randomly-selected user sync tones will fall into any of the bins at the satellite. Note that if we are taking our M^* bins from within the user channel itself, this actually has a beneficial effect in that it significantly reduces the probability of false detection P_{fd} as we get within about $M/2$ bins of proper alignment. In that case, about half of the total time,

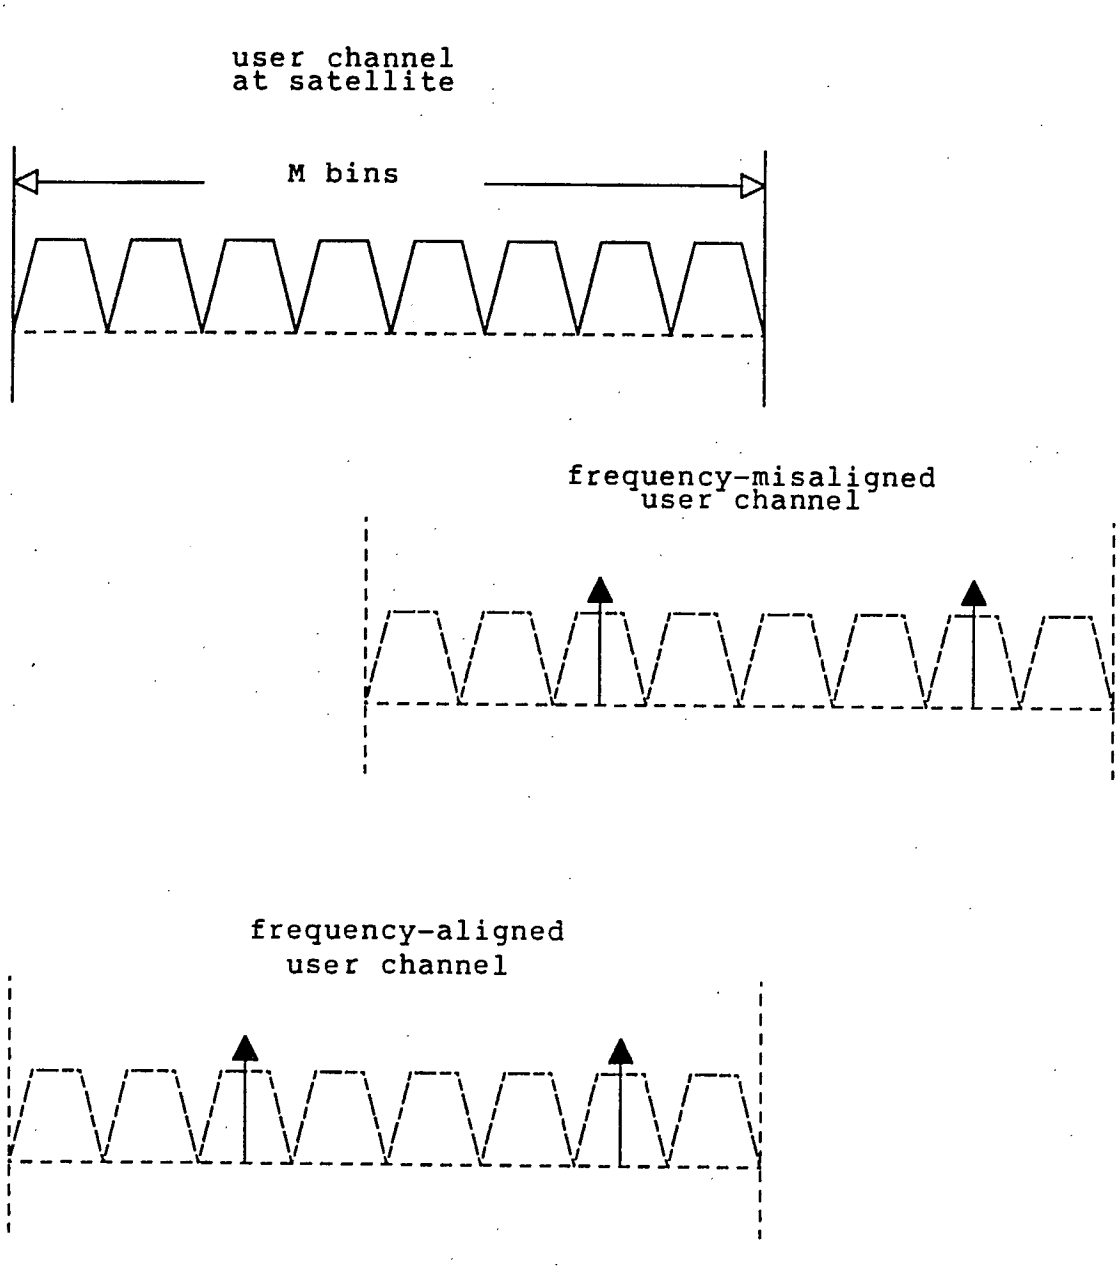

Figure 45: Frequency alignment errors with best 1-of-$M^* = M$ approach

the transmitted sync tones appear in the incorrect bins where they have a high probability of winning the best 1-of-M^* competition. The probability of seeing the intended tone bin win out is correspondingly reduced, lowering P_{fd}.

It is also clear from Figure 45 why we could not use a different approach based on looking for a consistent winner among any of the M^* bins in the hope of reducing the frequency search time. That would work for large frequency errors (such as depicted in the middle of the figure) only if the location of the sync tone probe was fixed in a particular bin. But as discussed earlier, this should be avoided because of susceptibility to MT-jamming.

Consider dedicating a portion of the user group resources for sync purposes. A frequency band of bandwidth B may be located at the edge of the FDMA user group as shown in Figure 46. This band is shared on a random access arrangement by any user wishing to synchronize. To keep the overhead acceptable, the sync band might be chosen to be, say, 10% of the user group bandwidth. With 1000 user channels for data transmission, this would create 100 noname "sync" channels. The SAW receiver will produce energy detector ouputs for each of the B/f_s tone bins in this band, just as it does for the channels in the active message portion of the user band. A best 1-of-M detection can be applied to these $S = B/Mf_s$ noname sync channels, and these S decisions can be supplied as data in a common sync information slot in the downlink.

A user attempting to synchronize starts by aiming for the center of the sync band, and transmitting a randomly selected tone in this middle channel. The user dwells here for a number of hops, then randomly selects a new channel to aim for (from the $S/2$ channels centred in the sync band), and transmits again. This is repeated over a number of random jumps within the sync band. We specifically aim for the middle $S/2$ channels so that when aligned with an error of no more than plus or minus $B/4$, all random jumps are sure to fall inside the sync band. Note that the user sync probe tone will have alignment errors with the bins detected by the SAW, so the procedure should be modified to transmit the probe signals at each of several alignments centered on the newly selected random position in the sync band. These will be separated by a ΔF chosen to ensure an acceptable amount of energy loss in the event of worst-case alignment $\Delta F/2$.

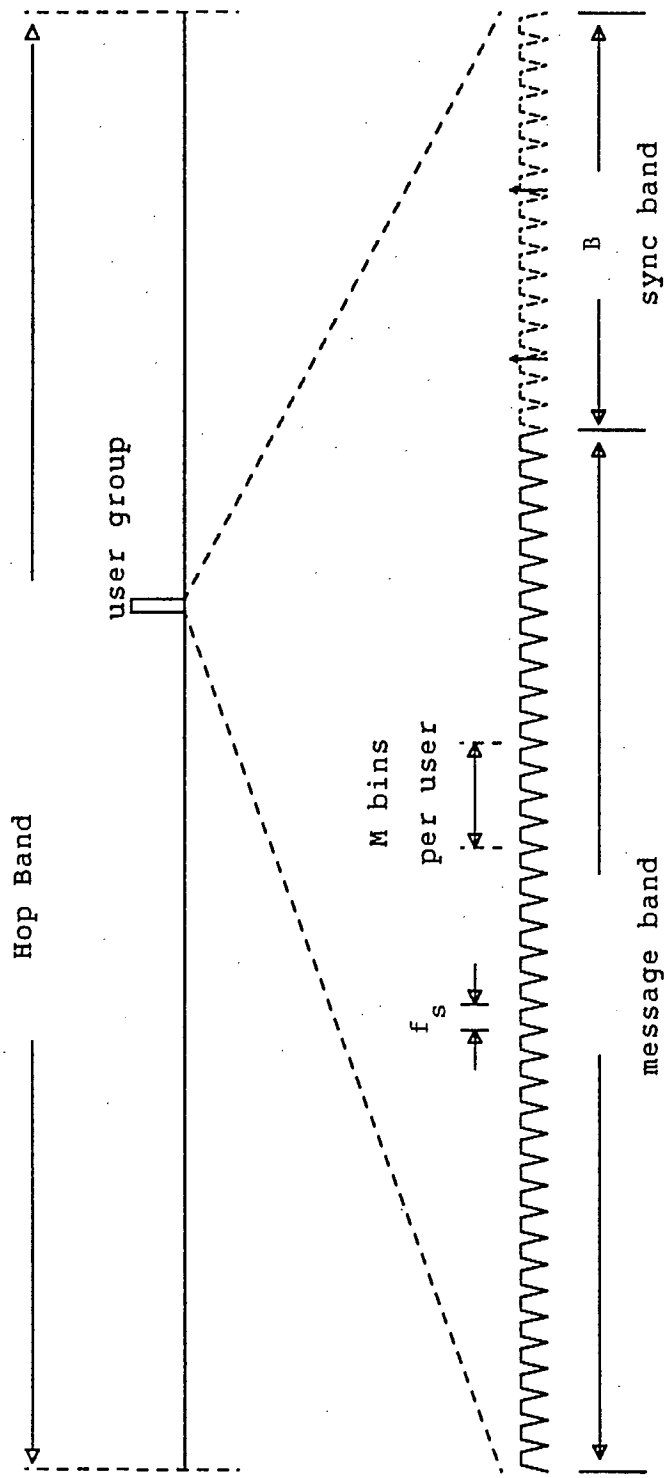

Figure 46: Sync-band approach

The downlink return is monitored for tone detections that are consistent with the random relative motions the user has selected within the sync band. The user is not concerned that the satellite is dividing the sync band up into M-ary channels for its own detection purposes, and is not concerned about which particular tone bin he happens to land in. He is only concerned about hitting one of these bins with a reasonably good alignment.

Note that the downlink return of the sync band channels will contain mostly random tone detections due to background noise and jamming, as well as some consistent detections due to other users who are attempting synchronization at that same moment. But each user will have his own unique randomly-selected jump pattern within the sync band, and so they will be clearly differentiable from one another, and from the random background detections. In addition, since synchronization will be a reasonably rare event, only a small fraction of the user population will be expected to be using the sync band at any one time.

If after a sufficient amount of observation time the user does not see the consistent detection pattern he is searching for, he steps his nominal carrier centre frequency by $B/2$ Hz and repeats the procedure. This process is of course repeated over all possible hopping sequence phases to be searched. Once the consistent pattern is seen, the user immediately knows his carrier frequency error to an accuracy of plus or minus f_s, and adjusts his carrier relative offset to aim for his normal user uplink channel. In this way, large initial carrier frequency alignment errors can be quickly reduced using these large $B/2$ frequency steps. This contrasts to the situation without a sync band in which users can only use their own normal uplink channels with a frequency step size ΔF chosen to yield an acceptable loss due to alignment error when the best search cell is probed (ΔF will be some fraction of the tone bin separation f_s).

The steps of sync-band approach are summarized as follows:

1. Aim for the center of the sync band and transmit a sync tone probe for a number of hops. Make small alignment adjustments (plus or minus $k\Delta F$, k a small integer) and repeat the transmissions at these alignments.

2. Aim for a new randomly selected location in the central 50% of the sync band, and

repeat 1). Repeat this step several times.

3. Monitor the downlink data returned for all channels in the sync band looking for relative motion of a detected tone that is consistent with the random motions selected in step 2). This is done with $(B/2)/f_s$ parallel hardware match detectors.

4. If no consistent pattern is seen, step the carrier frequency by half the size of the sync band, $B/2$, and go to 1).

5. If still no consistent pattern, step the hopping sequence phase by ΔT and go to 1).

6. Consistent pattern is seen; jump relative by your discovered frequency error to get to your normal uplink channel, and proceed to the verification phase.

Note that the procedure is amenable to pipelining as in Figure 7. This sync band procedure clearly has the possiblity of greatly reducing search time in the face of large initial frequency offsets. The reduction factor will be at most $(B/2)/Pf_s$ where P is a factor that accounts for i) the extra observation time needed because of the random jumping within the sync band, and ii) the increased observation time needed to make the detection of a consistent pattern as compared to simply looking for a detection in one particular position (as is done when searching without a sync band).

Note that we can think of this reduction as being due to the ability to process many different possible frequency alignments simultaneously, a form of parallel processing. This is exactly what is being done in trying to match up the tone detections evident in the downlink sync-band data with the relative motions imposed on the sync probe. We will have to check all $(B/2)/f_s$ possible start positions simultaneously, searching for consistent relative motion of a detected tone. This is made possible by having access to the detections for all B/f_s bins provided in a common portion of the downlink frame. This does have implications for the downlink overhead created by adoption of this sync-band approach. This common slot must be repeated in each of the downlink beam zones (see Figure 2). If each beam zone has only a fraction of the user population, this sync band data may represent a large overhead. One way to reduce the redundancy is to partition the sync-band into subbands, one for each user zone. Users in a particular zone use their own sync subband; the subband location and

size can be specified in the common downlink slot. This can reduce the downlink overhead by a factor of N_z, where N_z is the number of downlink beam zones. This means, however, that the frequency adjustment steps must be reduced in size to match the reduced size of the subbands, and the sync time will be increased. This is simply a consequence of reduced parallelism. Even if we are forced to make the size of the sync band (or subband) available to a user small enough to reduce the downlink overhead, we can still achieve useful reductions in the time-to- acquisition in the face of large initial carrier frequency errors.

Note that we cannot apply an analogous procedure in the hopes of reducing search time over the region of uncertainty in hopping sequence phases. That would require more than one frequency dehopper on the satellite.