Institut für Automation
Abt. für Automatisierungssysteme

Technische Universität Wien

Projektbericht Nr. 183/1-121
June 2002

On the Problem of Internal Interference in CDMA-Based Ad-Hoc Networks

Günther Gridling

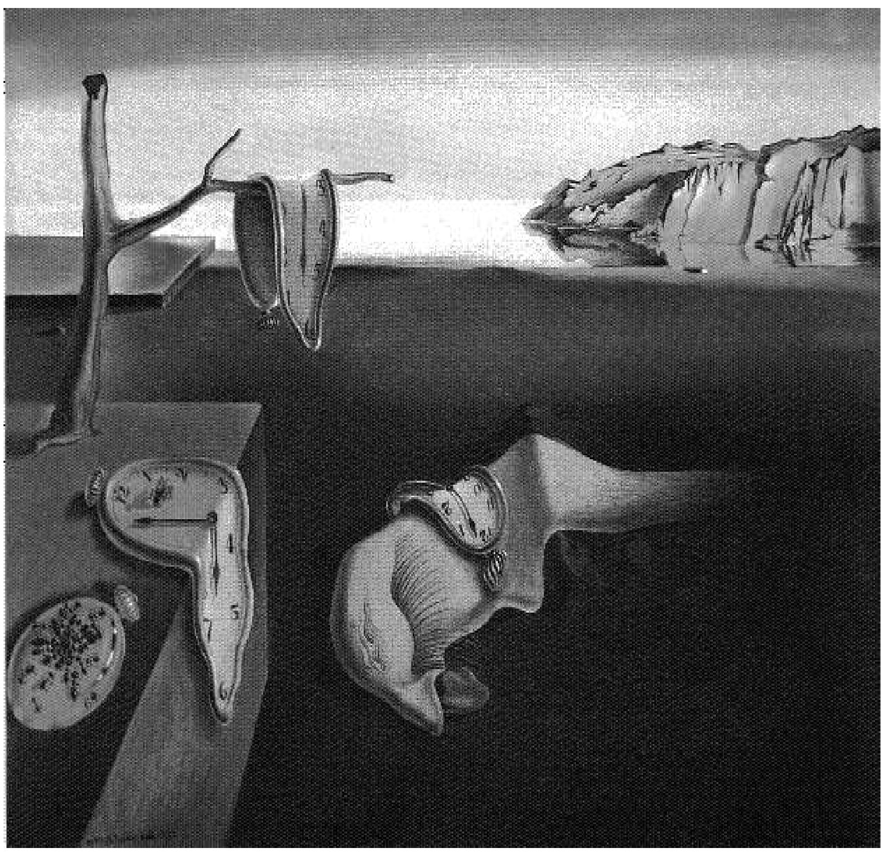

Salvador Dali, "Die Beständigkeit der Erinnerung"

On the Problem of Internal Interference in CDMA-Based Ad-Hoc Networks

Günther Gridling

July 17, 2002

Abstract

This report briefly examines the feasibility of CDMA-based communication in wireless ad-hoc networks. In particular, it deals with the problem of internal interference, i.e., the interference one communication channel suffers from the concurrent activity on channels in its vicinity. A simulation framework and some simulation results are presented which suggest that with some loose constraints on the topology, the interference problem could be handled with reasonable expenditure of hardware resources.

1 Motivation

The W2F-project[1] aims at developing a next-generation fieldbus operable on wireline as well as wireless networks. Currently, the designated communication technique is (UWB-)CDMA. While CDMA offers significant advantages in terms of robustness and security, it also poses certain problems, especially in conjunction with the ad-hoc nature of a W2F network. Virtually all current CDMA-based systems are tailored to the use of base stations, allowing for the use of dedicated pilot channels to synchronize different communication channels. Obviously, on an ad-hoc network, such synchronization is not possible. In the course of outlining the basic structure of the W2F network, the question came up whether the use of completely asynchronous CDMA in a wireless network —where base stations and shared control channels are not an option— is at all feasible. To investigate the problem, a simple simulation environment was developed, which will be described below.

1.1 W2F Network Structure

Since the simulation is based on the W2F network as currently envisioned, we should first list the most basic features of such a network:

1. The number of nodes is assumed to be in the range of 100 up to 1000.
2. Each node can communicate over at least k channels concurrently, with k in the $3-10$ range.
3. For each communication channel, a separate CDMA key is used (?).

[1] This research is part of our W2F-project, which targets a wireline/wireless fieldbus based upon spread-spectrum (CDMA) communications, see *http://www.auto.tuwien.ac.at/Projects/W2F/* for details. W2F is supported by the Austrian START programme Y41-MAT.

4. CDMA key length is projected to be about 1000 chips to achieve a sufficient processing gain (see [Goi98] for details).
5. To facilitate reasonably accurate power control, communication should be bidirectional, but it may turn out to be unidirectional if absolutely necessary.
6. Single points of failure are to be avoided, which precludes the concept of base stations and pilot channels.

1.2 CDMA Communication

In the following, we will try to give a brief description of the principle of CDMA communication(again, see [Goi98] for details).

CDMA communication relies on the use of a different (usually binary) code sequence per communication channel, where each data bit is encoded with one complete key sequence in the following manner:

A logical *zero* is mapped to a numerical value of $+1$, and a logical *one* to a value of -1. This transformation is applied both to the data bits and the bits of the key sequence, the so-called *chips*. To encode a bit, a complete transformed key sequence is multiplied with the numerical value corresponding to the data bit. After that, the signal representing the particular data bit is constructed by multiplying the sequence of ± 1-values with an appropriate transmission level.

The receiver multiplies its incoming signal with the sequence of ± 1-values representing the key and integrates over a complete sequence. Based on the integral value, it then decides whether the sender has transmitted a logical one or zero. With perfect codes and absent *external* interference (any interference not produced by the network itself), the absolute value of the integral would always be $numberofchips \cdot signallevel$, but it would be positive for a zero bit, and negative for a one bit. The trick is, of course, that due to having opposite sign on exactly half their chips, transmissions using other codes are canceled out through the integration. This, however, only works for perfect codes. Two codes out of an imperfect code set do not always differ in exactly half the chips, so different codes are not canceled completely, thus producing what we will call *internal* interference. The number of chips that are not canceled between two codes is called *cross-correlation* of these codes.

Also, if the codes are balanced, i.e., the number of ones is equal to the number of zeroes, any DC components in the external interference are eliminated in the same way.

2 The Interference Problem

One of the major problems in CDMA-based communication is the lack of perfect codes. While such codes are available for synchronous systems (e.g., Walsh codes), none are known that work in asynchronous systems. With imperfect codes, however, the problem of internal interference is introduced. Its impact is all the more severe in an ad-hoc network such as W2F, since the distances of nodes are highly arbitrary, thus widening the impact of the so-called near-far-problem: It must be considered that any wireless transmission is attenuated with the square of the distance, in indoor-environments even more. When a node listens to a remote sender, even a small cross-correlation with another channel can severely impair communication if the other channel's sender is much nearer and its signal is received at a correspondingly higher level.

Although the W2F network seems rather unsuited to CDMA communication (primarily due to its ad-hoc topology), it offers one decisive advantage over base station systems: The transmission range is restricted to a minimum. In common base station systems, any mobile node p usually has to use a considerable transmission power in order to reach the nearest base station. A large number of other mobiles can be expected to be within p's transmission range, and many of them will be much nearer to p than the base station. Therefore, uplink and downlink usually employ different frequency bands, effectively limiting the maximum bandwidth for both. Without such a dual-band approach, the nodes will of course suffer a considerable and often intolerable interference from this transmission. This, in turn, increases the demands on the hardware deployed for multi-user interference reduction considerably. In a W2F network, however, the transmission range only needs to be large enough for a node to have k other nodes in its range, so the transmission power and hence the overall internal interference will usually be substantially smaller. The question remains, however, whether this improvement in signal to noise ratio (SNR) due to the decrease in internally produced interference is sufficiently large to bring down the MUIR-related expenditure of hard- and software to an acceptable level.

2.1 Internal Interference

(Note: This chapter was originally intended to supply a sound framework for the ring model to be detailed in Chapter 3. However, currently it is neither complete nor consistent. Originally, the model was no more than an intuitive means of argumentation that the problem of internal interference is not necessarily fatal. To establish a proper technical context, some work needs to be done. Since the author's dissertation has nothing whatsoever to do with CDMA communication in particular, this work will be postponed for the time being.)

2.1.1 Definitions

Definition 1 (Interference Function) *Assume two nodes s and r. The function*

$$f_I(t, s \rightarrow r) = I^{(tot)}(t, s)r \tag{1}$$

gives the total interference the link $s \rightarrow r$ suffers at time t.

Definition 2 (Cross Correlation Function) *Assume two pairs of nodes $s_i, r_i, 1 \leq i \leq 2$. The function*

$$f_X(t, C^{(s_1 \rightarrow r_1)}, C^{(s_2 \rightarrow r_2)}) = X^{(s_1 \rightarrow r_1)}_{(s_2 \rightarrow r_2)}(t) \tag{2}$$

gives the cross correlation between two CDMA transmissions $s_1 \rightarrow r_1$ and $s_2 \rightarrow r_2$ at time t.

Definition 3 (Code Replacement Latency) *The* code replacement latency T_L *is the minimum time that a code will be in use under normal conditions. I.e., the minimum time it takes the system to* specifically *change a code in response to interference problems. In particular, this is the time frame which has to be accounted for in any analysis of interference behavior.*

Definition 4 (Interference Granularity) *The* interference granularity G_I *is the maximum value an interference may change while still being considered constant for the interference analysis.*

Definition 5 (Constant Interference) *The interference on a link is called* constant *during a time period* $t_{start} - t_{end}$*, if the change of the corresponding interference function within any duration of one code latency is bounded by the interference granularity:*

$$\forall t_1, t_2 \in [t_{start}, t_{end}], t_1 < t_2, t_2 - t_1 \leq T_L : |f_I(t_1, s \rightarrow r) - f_I(t_2, s \rightarrow r)| \leq G_I \tag{3}$$

Definition 6 (Interference Density) *The* internal interference density *is the level-weighted sum of all signals from all nodes that arrive at a particular point:*

2.1.2 Types of Internal Interference

When we speak of internal interference on a communications link based on asynchronous CDMA, we will distinguish between three types of interference: *unsigned*, *signed* and *balanced*.

Definition 7 (Unsigned Interference) *The interference on a link is called* unsigned *during a time period* $t_{start} - t_{end}$*, if the corresponding interference function has a positive maximum and a negative minimum within this time period.*

$$\forall t \in [t_{start}, t_{end}] : max(f_I(t, s \rightarrow r)) > 0 \wedge min(f_I(t_2, s \rightarrow r)) < 0 \tag{4}$$

Unsigned interference is called symmetrical*, if the absolute values of maximum and minimum that can occur are equal. Correspondingly, if the absolute values are different, it is called* asymmetrical.

Definition 8 (Signed Interference) *The interference on a link is called* signed *during a time period* $t_{start} - t_{end}$*, if within this period, maximum and minimum of the corresponding interference function have equal signs, or one of them is zero.*

$$\begin{aligned}\forall t \in [t_{start}, t_{end}] \quad : \quad & sgn(max(f_I(t, s \rightarrow r))) = sgn(min(f_I(t_2, s \rightarrow r))) \\ & \vee\, max(f_I(t, s \rightarrow r)) = 0 \\ & \vee\, min(f_I(t_2, s \rightarrow r)) = 0 \end{aligned} \tag{5}$$

Definition 9 (O_I-Balanced Interference) *A signed interference on a link is called* balanced*, if the distribution of the possible values of the corresponding interference function is symmetrical around zero with a small offset. This offset is called the* interference imbalance O_I *and is given in percent of the perfect signal level for that link. When* O_I *is zero, we speak of* perfectly balanced interference.

Signed and balanced interference usually originate from transmissions which, while not really synchronous, at least use the same number of chips per bit and chip duration as the link they interfere with, henceforth referred to as "count/duration-synchronous" (c/d-synchronous). This type of interference produces a deterministic, signed interference value depending on the cross-correlation between the respective codes and the distance. This means that a strong interference from one particular transmission can actually have a favorable effect on a node's SNR in that it may counter-balance other interferences with reverse sign. The cross correlation properties of the underlying CDMA codes determine whether the interference is balanced or merely signed.

Completely asynchronous interference, on the other hand, will not be deterministic, hence it must be handled as a worst-case value, which is, of course, unsigned (hence "unsigned" interference). Therefore, no counter-balancing will occur, leading to much higher total interferences compared to the signed and more so the balanced model. Note that any external interference will usually have to be regarded as completely asynchronous and hence unsigned.

2.1.3 Internal Interference in an Ad-hoc Network

It is rather obvious that the design of a wireless CDMA-based ad-hoc network will have to rely on the internal interference being mostly O_I-balanced (with the value of O_I depending on the interference density), since with realistic cross-correlation values and disregarding counter-balancing effects, interference soon exceeds acceptable limits. Note that, even in a model with perfectly balanced interference, a worst-case scenario could occur where all codes in a certain area have high correlation values with equal polarity. This means that in the worst case, the perfectly balanced interference model could degrade to the unsigned one. This, however, could be —and would have to be— circumvented, e.g., by randomly changing codes in areas with high internal interference or by assigning codes through a location-aware service which optimizes interference through code distribution.

Assuming internal interference to be balanced, however, is not sufficient to bring the overall interference down to an acceptable level. In a network with number of nodes $n = 200$, redundancy $k = 3$, maximum cross correlation $C^{(max)} = 10\%$, minimum cross correlation $C^{(min)} = -10\%$, attenuation $a = 0$ (wireline) and decision threshold $Th_D = 30\%$, each node receives all 600 transmissions with the same signal level. With the probabilities $P(C^{(max)}) = P(C^{(min)})$, we can expect 300 transmissions per correlation value. However, an imbalance of just two transmissions would be sufficient to raise the interference above the decision threshold, and the probability of that occurrence ist clearly too high.

3 "Ring" Model

To roughly estimate the scale of the interference problem in a wireless W2F-network, we tried to abstract the influence of attenuation with the following simple —and rather artificial— model:

Assume a node p transmitting a CDMA-based signal to nodes n_1 and n_2, which are in distances $d(p, n_1)$ and $d(p, n_2)$ from p, $d(p, n_1) < d(p, n_2)$. Any transmission is attenuated with attenuation a, so if p transmits a signal at level $L^{(snd)}$, it is received at n_i with level

$$L^{(rcv)} = A(d(p, n_i), L^{(snd)}) = \frac{L^{(snd)}}{d(p, n_i)^a}. \tag{6}$$

The codes $C^{(p \to n_1)}$ and $C^{(p \to n_2)}$ used for these transmissions have a cross-correlation of $X(C^{(p \to n_1)}, C^{(p \to n_2)})$, or $X^{(p \to n_1)}_{(p \to n_2)}$ in short. Each of the links $p \to n_i$ suffers some unipolar interference $I^{(tot)}_u(p \to n_i)$. The interference $I^{p \to n_2}(p \to n_1)$ which node n_2's receiver link to p suffers from node p's transmission to n_1 is

$$I^{p\to n_2}(p\to n_1) = X^{(p\to n_2)}_{(p\to n_1)} \cdot A(d(n_1,n_2), L^{(min)}(n_1)) = X^{(p\to n_2)}_{(p\to n_1)} \cdot \frac{L^{(min)}(n_1)}{d(n_1,n_2)^a}. \tag{7}$$

Likewise, $I^{p\to n_1}(p\to n_2)$, the interference on $p\to n_1$ due to the transmission $p\to n_2$, would be

$$I^{p\to n_1}(p\to n_2) = X^{(p\to n_1)}_{(p\to n_2)} \cdot A(d(n_1,n_2), L^{(min)}(n_2)) = X^{(p\to n_1)}_{(p\to n_2)} \cdot \frac{L^{(min)}(n_2)}{d(n_1,n_2)^a}. \tag{8}$$

For the sake of simplicity, assume that there are no other nodes present, and that $I_u^{(tot)}(p\to n_1)$ is constant at all points within distance $d(p,n_2)$ from p. In that case, the total interference which $p\to n_1$ suffers is equal to the unipolar interference plus the interference from $p\to n_2$:

$$I^{(tot)}(n_1) = I_u^{(tot)}(p\to n_1) + |I^{p\to n_1}(p\to n_2)|. \tag{9}$$

If the threshold for the bit-decision is Th_D, $p\to n_1$ has some "interference slack" for $|I^{p\to n_1}(p\to n_2)|$, namely

$$I^{S(p\to n_2)}(p\to n_1) = Th_D - I_u^{(tot)}(p\to n_1). \tag{10}$$

Obviously, $|I^{p\to n_1}(p\to n_2)|$ will increase with decreasing $d(p,n_1)$, which raises the following question: with a given $d(p,n_2)$, what is the minimum $d(p,n_1)$ required so that n_1 can tolerate $I^{p\to n_1}(p\to n_2)$?

Assuming that $I^{S(p\to n_2)}(p\to n_1)$ is positive (otherwise, the unipolar interference would already be intolerably high), we can give the requested minimum $d(p,n_1)$ as

$$d_{min}(p,n_1) = d(p,n_2)\cdot \sqrt[a]{\frac{I_u^{(tot)}(p\to n_1)}{I^{S(p\to n_2)}(p\to n_1)}}. \tag{11}$$

Thus, from the distance between p and its farthest peer n_2 and the maximum tolerable proximity, we could define an "interference tolerance ring" around p for the position of n_1. As long as n_1 is located within the inner and outer border of this ring, the interference it suffers from $p\to n_1$ is tolerable, and no MUIR is needed. This is, of course, an oversimplified model, since in real systems, the interference is more complex. The "ring" will usually be an irregular structure, and because interference from other nodes is signed, even isolated areas of interference tolerance could occur. However, the principal effect should still be noticeable even in a realistic environment.

With this in mind, we developed a simple simulation in which a number of nodes were placed randomly within a rectangular area. The distance between any two nodes was not allowed to fall short of some d_{min}, to avoid two nodes being so close as to render all communication to other nodes impossible. As a first (and optimistic) approach, we chose to ignore the carrier detection problem and set the bit decision threshold to 0. The maximum interference I_{max} was assumed to be 70 percent.

The results looked promising: With an (admittedly extremely optimistic) C_{max} of 3 percent, in most cases no node violated the ring requirement.

4 Simulation

Encouraged by these results, we developed the simulation further to handle the interferences more realistically. For a given number of nodes in a given (rectangular) area, the simulation now computes the transmission level produced by each node and sums up the effect on all other nodes to determine the interference each node suffers. In the following, we will briefly describe the way this computation is done.

4.1 Interference Computation

Let $n_1(p)$ ... $n_k(p)$ be the k neighbors node p talks to (with (p) omitted when unambiguous), in distances $d(p, n_1)...d(p, n_k)$ from p, respectively. Each node p is capable of transmitting a signal at the maximum level $L_p^{(max)}$. For a node q to be able to correctly correlate a signal from node p, the receive level $L_p^{(rcv)}(q)$ of the signal must be at least $L_q^{(min)}$. Depending on the distance $d(p, q)$, the attenuation exponent a attenuates a level $L_p^{(snd)}$ sent from node p to node q to

$$L_p^{(rcv)}(q) = A_{d(p,q)}(L_p^{(snd)}) = \frac{L_p^{(snd)}}{d(p,q)^a}. \tag{12}$$

So, to talk to neighbor $n_x(p)$, node p uses an optimal transmission level

$$L_p^{(opt)}(n_x) > L^{(min)} \cdot d(p, n_x)^a. \tag{13}$$

With non-perfect power control, the transmission level will usually be higher than that and might occasionally fall below, causing transmission failures. Obviously, if $L_p^{(opt)}(q) > L_p^{(max)}$, q is out of p's range, but any transmission from p still produces interference at q which must be considered, however small it may be.

So, assuming perfect power control, the signal level a node p produces is

$$L^{(sum)}(p) = \sum_{i=1}^{k} L_p^{(opt)}(n_i). \tag{14}$$

However, due to the special features of CDMA, the interference some other node q suffers from this signal is not simply proportional to the quotient $L^{(sum)}(p)/d(p, q)$. Rather, each link must be considered separately to account for the different cross correlations. E.g., given two links with a cross correlation of zero between the respective CDMA keys, a transmission over one link will not produce an interference on the other link, regardless of transmission level and distance. On the other hand, with sufficiently high cross correlation, $L^{(rcv)}$ need not be high to produce unacceptable interference.

Therefore, in the simulation, the computation procedure is as follows:

for each node p
  for each neighbor $n_i (1 <= i <= k)$:
    compute $L^p_{snd}(n_i) = L^p_{opt}(n_i) = L_{min} \cdot d^a_{p,n_i}$
    for each node q
      compute $L^q_{rcv}(p) = A^{d_{p,q}}(L^p_{snd}(n_i))$
      add $L^q_{rcv}(p)$ to the interference sum of node q

4.2 Simulation Parameters

The following parameters can be adjusted interactively:

number of nodes n**:** Due to the fact that the simulation was written from scratch in C++, as opposed to using a simulation environment like Matlab, it is pretty fast. Hence, the number of nodes can be in the 1000s, while the duration of one simulation run is still acceptable (see Table 4.2).

required redundancy k**:** Can be any value between 1 and n. Of course, with exceedingly high values, the internal interference grows to fatal values.

maximum transmission signal level L_{max}**:** See description of L_{min} below.

minimum required receive signal level L_{min}**:** Maximum transmission level and minimum receive level (together with the attenuation) basically just determine the maximum transmission range any node can achieve. The actual transmission range depends on k and will usually be a lot smaller than the maximum.

spatial distribution $x_{range} \cdot y_{range} \cdot z_{range}$**:** Nodes are placed randomly within a cuboid with the given dimensions.

minimum distance d_{min}**:** Whenever the distance between two nodes is too small, the near-far problem becomes dominant. To avoid this, a minimum distance should be observed by the network designer. Note that with increasing minimum distances, it is not always possible to position all nodes accordingly, so the simulation program only gives a best effort: for each node, random positions are chosen and checked for minimum distance with all other nodes. If by the tenth attempt no valid position is found, the node will be placed at the position with the largest minimum distance, even if it falls short of the user-specified minimum distance.

cross interference range $[C_{min}, C_{max}]$**:** Actual interference values for each pair of codes are randomly chosen from the interval $[C_{min}, C_{max}]$. Note that this is not an accurate model of cross correlations among codes of usual code sets. Gold-codes, for example, can only have three different cross correlation values, and those are slightly imbalanced. However, the results of the simulation are only used as an estimation of real-life behavior anyway, so the error introduced should be acceptable.

interference threshold: This is the maximum interference any node can tolerate. It depends mostly on the amount of external interference that is to be expected as well as whether there are carrier detection mechanisms other than the correlator output.

attenuation a **or attenuation range** $[a_{min}, a_{max}]$**:** If a single value is given for the attenuation, it is assumed to be globally constant and is applied to any link. If an interval is given, the actual attenuation for each link is randomly selected from that interval.

bidirectional or unidirectional: While the network can be required to be bidirectional, this does not mean that the program will optimize the links with regard to transmission ranges and the amount of internal interference produced. Rather, nodes are iterated in no particular order, and each node chooses its nearest neighbors as peers, if possible. If it has already been chosen as neighbor by one or more other nodes, it must in turn choose those nodes as neighbors, even if there are other nodes much closer. Therefore, unnecessary long transmission ranges may occur. We chose to not optimize to account for a different problem with the network topology: With randomly placed

number of nodes	redundancy	execution time (ms)
50	3	< 10
100	3	30
150	3	90
200	3	170
250	3	280
300	3	430
50	5	< 10
100	5	50
150	5	130
200	5	230
250	5	390
300	5	580

Table 1: Execution Times

nodes, network partitions might occur. While each node might have k peers, there could be a group of at least k nodes none of which has a connection to a node outside the group, effectively separating this group from the network. To overcome this problem, some nodes might have to extend their transmission range, counter-balancing the gain from optimization of bidirectional links. As long as no partitioning detection is implemented, we chose to leave bidirectional connections unoptimized, so at least in a bidirectional network, transmission ranges should be realistic.

4.3 Simulation Output Values

From the parameters described above, the program creates a network and computes the following values:

number of nodes without connection n_{sep}: This can be used as a reality check for the parameters $(x \cdot y \cdot z)$, L_{max}, L_{min}, and a. Whenever this number is not zero, those parameters have to be adjusted accordingly.

number of nodes with a connectivity below the required redundancy $n_{<k}$: Like n_{sep}, this value must be observed and the simulation parameters set accordingly.

average number of peers per node n_{peers}: Depending on the maximum transmission range, this value gives the average number of nodes one node could possibly communicate with.

maximum interference among all nodes C_{max}: See description of C_{avg} below.

average interference C_{avg}: The above values are closely related. For each node, the interference each link suffers from all ongoing transmissions is computed. C_{max} gives the maximum value that occurred at any link for any node, C_{avg} gives the average.

number of nodes that suffer an interference above the threshold $n_{>th}$: This is the number of nodes that suffer an internal interference in excess of the maximum acceptable value on at least one link. The affected links will not be operable most of the time.

number of nodes that need at most 10 MUIRs: If the interference threshold is exceeded at some link, a node must perform multi-user interference reduction (MUIR) for this link to be useful. While for many erratic links, it will be sufficient to do MUIR for the strongest interferer, some may need more MUIRs. The simulation computes the number of nodes that can manage with 1 ... 10 MUIRs.

number of nodes for which 10 MUIRs are still insufficient: These are the nodes that have at least one link which cannot be saved even by eliminating the 10 strongest interferers.

4.4 Simulation Output Graphs

Furthermore, a number of graphs can be generated:

1. number of nodes without connection
2. number of nodes with less than k connections
3. number of nodes that suffer an interference above the threshold
4. number of nodes that need at most 10 MUIRs
5. number of nodes which need more than 10 MUIRs

Those graphs are generated by running the simulation multiple times. In the simplest case, the simulation is run a user-specified number of times, each time using the same parameter set. Alternatively, the user can specify a start, end and increment value for the following parameters:

1. number of nodes
2. maximum transmission level
3. attenuation minimum
4. attenuation maximum
5. minimum distance
6. interference threshold
7. cross correlation minimum
8. cross correlation maximum
9. cross correlation absolute limit

5 Simulation Results

The reason this simulation was developed in the first place was that we needed to find out whether the problem of internal interference in a homogeneous ad-hoc network is indeed as overwhelming as it seemed. The most important result we gained from this simulation is that this seems not to be the case. While it is easy to mess up in the network design and cause internal interference to be fatal, it seems perfectly possible to carefully design a network so that the percentage of nodes requiring MUIR is in the single figures, especially when all or most of the nodes are immobile.

The other results mentioned here are rather straightforward, but still it is handy to see them in cold print. E.g., the minimum distance between any two nodes is, of course, a vital parameter in the network design. Obviously, as soon as nodes come too close, each of them will need to compensate for the other node's transmissions, so each of them will need at least k MUIRs. To tackle this problem, the following suggestion by W2F team member Hannes Stratil seems to be a workable solution: any nodes which are that close to each other should go into a dedicated cooperation mode in which they appear to other nodes as one node with more extensive hardware resources, scheduling their links to other nodes. That way, it could be possible for the network designer to produce largely arbitrary hardware redundancy out of a pool of identical W2F-nodes by simply positioning an appropriate number of them close together.

In a network topology with no restrictions, however, a lot of nodes are so close to each other as to produce interference on some channels, but not close enough so that the above mentioned cooperation mode would be viable. Therefore, it is vital that the nodes are as evenly distributed as possible. Depending on the restrictions regarding the minimum distance as well as the number of mobile nodes, each node must be equipped with a number of extra receivers to be used for MUIR. Typically, about three should do, but this obviously depends very much on the specific network properties. Note that those extra receivers need not be dead weight. In order to compensate for an interfering channel via MUIR, a node must know the code used on that channel. To make a virtue of necessity, this could be instrumented as an implicit multicast functionality. Security and redundancy allowing, static nodes could combine interfering channels to form multicast channels, thus reducing the internal interference in the area considerably.

The simulation also showed that in well-distributed topologies, even maximum cross correlations of 20 percent can be tolerated, which means that, e.g., Gold codes with a minimum m (see [Goi98]) of 7 (translating to a minimum code length of 129 chips) could be used.

In the following, a few plots of simulation runs are presented. The parameters used for the simulation are listed as follows: On top of the the plot, you can see whether the network was configured to be uni- or bidirectional. In the rest of line one, the number of nodes, the required redundancy, and the maximum transmission range are listed. Line two gives the transmission level (tx), the minimum required receive level (rx), the spatial distribution, and the minimum distance. The last line on top of the plot shows the minimum and maximum interference (the actual interference between any two codes is some random value in this range) as well as the interference threshold. If this threshold is exceeded, a node has to make use of MUIR. Below the plot, the resulting values are shown. First the percentage of nodes without connection, then the percentage of nodes with less than k connections, and as the last value in line one the average number of peers per node. Line two and three show the resulting internal interference. Line two gives the absolute maximum interference that occurred, the average of positive and negative interference, the number and percentage of nodes that exceed the interference threshold on at least one link, and finally the number and percentage of nodes that cannot manage even with ten MUIRs. In line three, you see the number and percentage of nodes that exceed the threshold, but can manage with less than ten MUIRs. The numbers are given for each MUIR count up to ten. The last line shows the average transmission range used.

As an example, Figure 1 shows a network of 100 nodes in an area of $100 \cdot 100$ units, $k = 3$, where no relevant minimum distance was observed. Compare that to Figure 2, which shows a network with equal parameters, but with a minimum distance of 5 units.

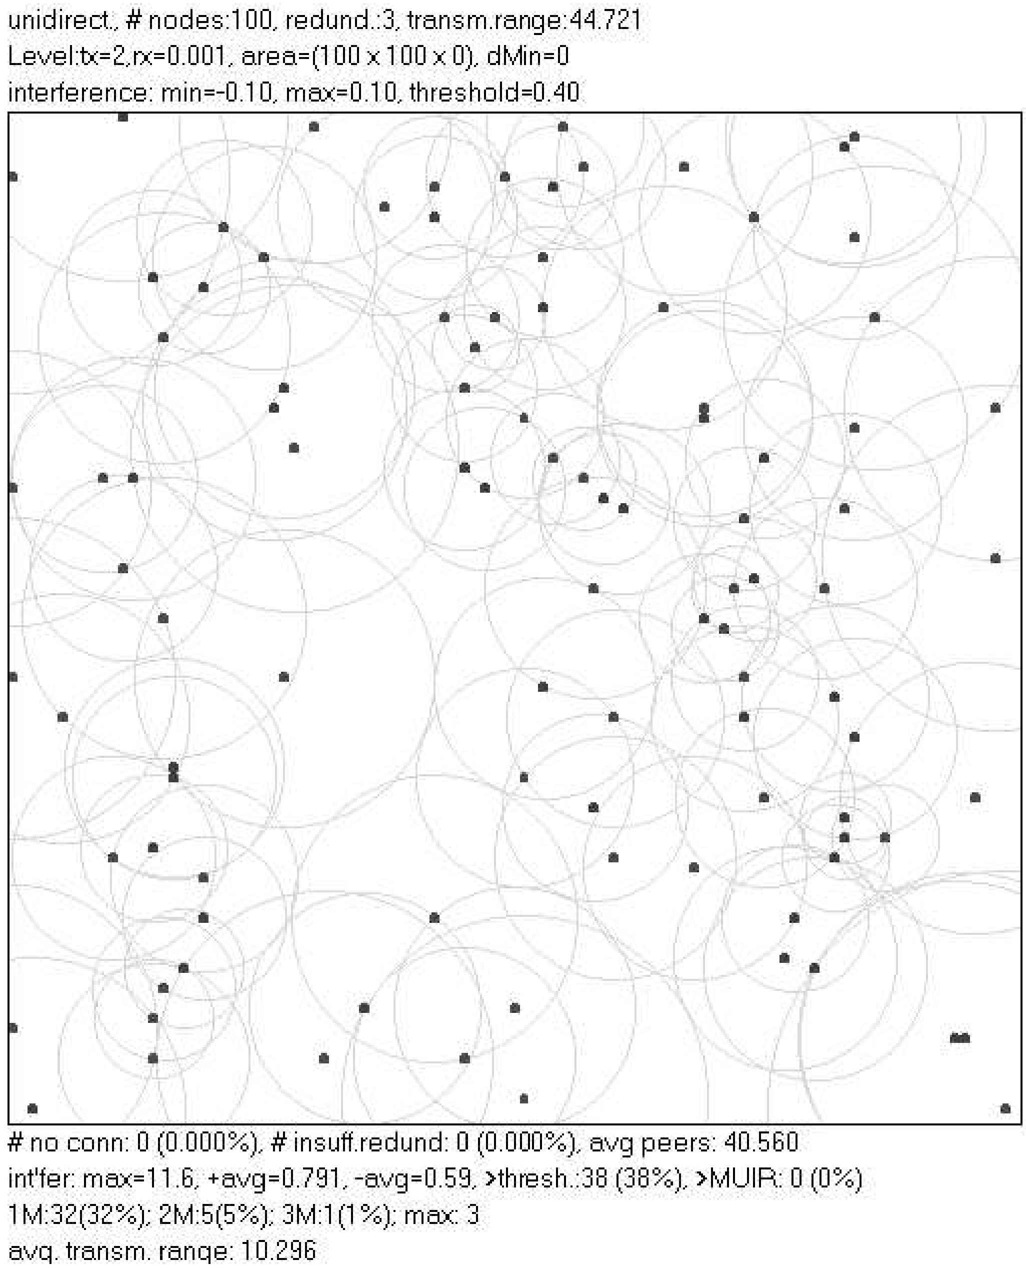

Figure 1: *No Minimum Distance Observed*

As mentioned above, the simulation can also produce graphs by varying some user-specified parameter, the so-called loop parameter. In that case, the user gives start value, end value, and step size for the loop parameter, along with a repeat count. The simulator assigns the start value to the loop parameter and creates a number of networks as specified with the repeat count value, all with the same parameter set. The resulting values are averaged and stored. Then, the step size is added to the loop parameter, and the procedure is repeated until the maximum value is exceeded. In the following figures, the repeat count is usually 10, unless specified otherwise.

In Figure 3 below, the influence of the minimum distance on the internal interference is shown. A minimum distance of $6 - 7$ units is optimal for this particular network. Note that after ten random placements of a node have failed to comply to the minimum distance requirement, this particular node is placed in violation of the requirement. As the minimum distance grows larger, it becomes increasingly difficult to place all nodes correctly, hence the internal interference increases.

With regard to the connectivity, the simulation showed that assuming an attenuation exponent of 2 and some appropriate minimum distance, the number of nodes necessary to ensure k-redundancy can easily be determined. Figure 4 shows a plot with varying node counts (repeat count was 30). Note that the graphs are scaled to their maximum. Unscaled, the same plot would look like Figure 5.

In Figure 6, the loop parameter is the interference threshold, running from 20 to 100 percent. A threshold of 20 percent is exceeded by about 42% of the nodes. A threshold of 85 percent is rarely exceeded.

The attenuation greatly influences the internal interference. In Figure 7, the attenuation maximum runs from 2 to 5, while the minimum is constant at 2. With higher attenuation, the number of links on which the interference threshold is exceeded comes close to 100 percent, but this plot shows that with as little as 3 MUIRs, most of the links can be saved.

In Figure 8, the attenuation minimum loops from 2 to 5, while the maximum is constant at 5. (Note that the maximum transmission level has been increased, since the average of the attenuation now runs from 3.5 to 5 instead of from 2 to 3.5.) This plot shows that it is not just a high attenuation which deteriorates the internal interference. Rather, the distribution of the attenuation is decisive. Still, with three MUIRs, most links can be saved.

Finally, in Figure 9, the absolute cross correlation between any two codes loops from 3 to 30 percent, with an attenuation exponent of 5. Naturally, with higher cross correlation values, the internal interference increases, but again, most links can be saved with 3 MUIRs.

References

[Goi98] Alois M.J. Goiser. *Handbuch der Spread-Spectrum Technik.* Springer, Wien, New York, 1998.

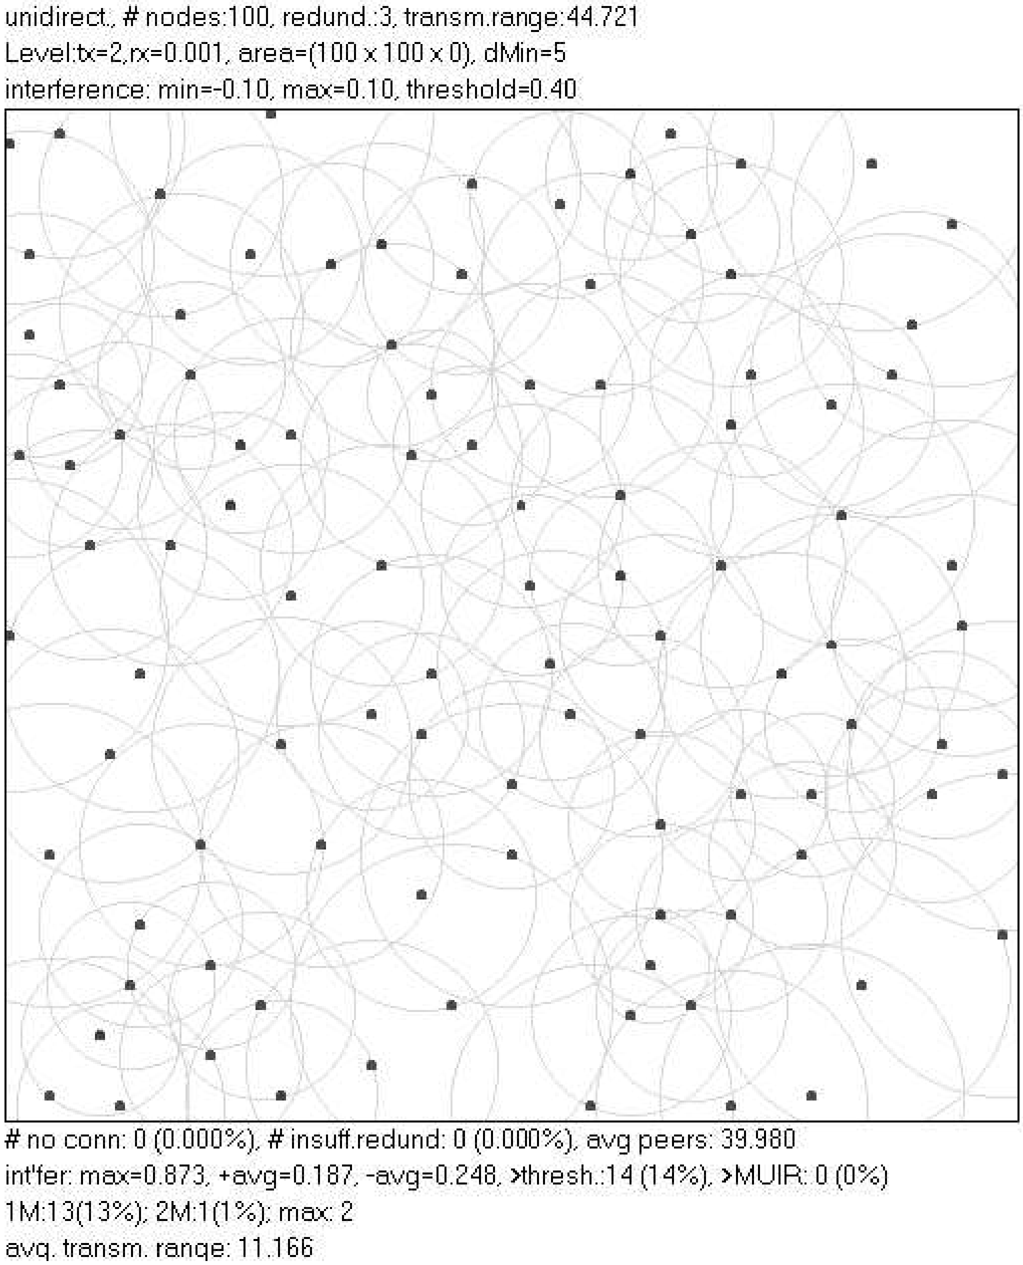

Figure 2: *Minimum Distance of* 5

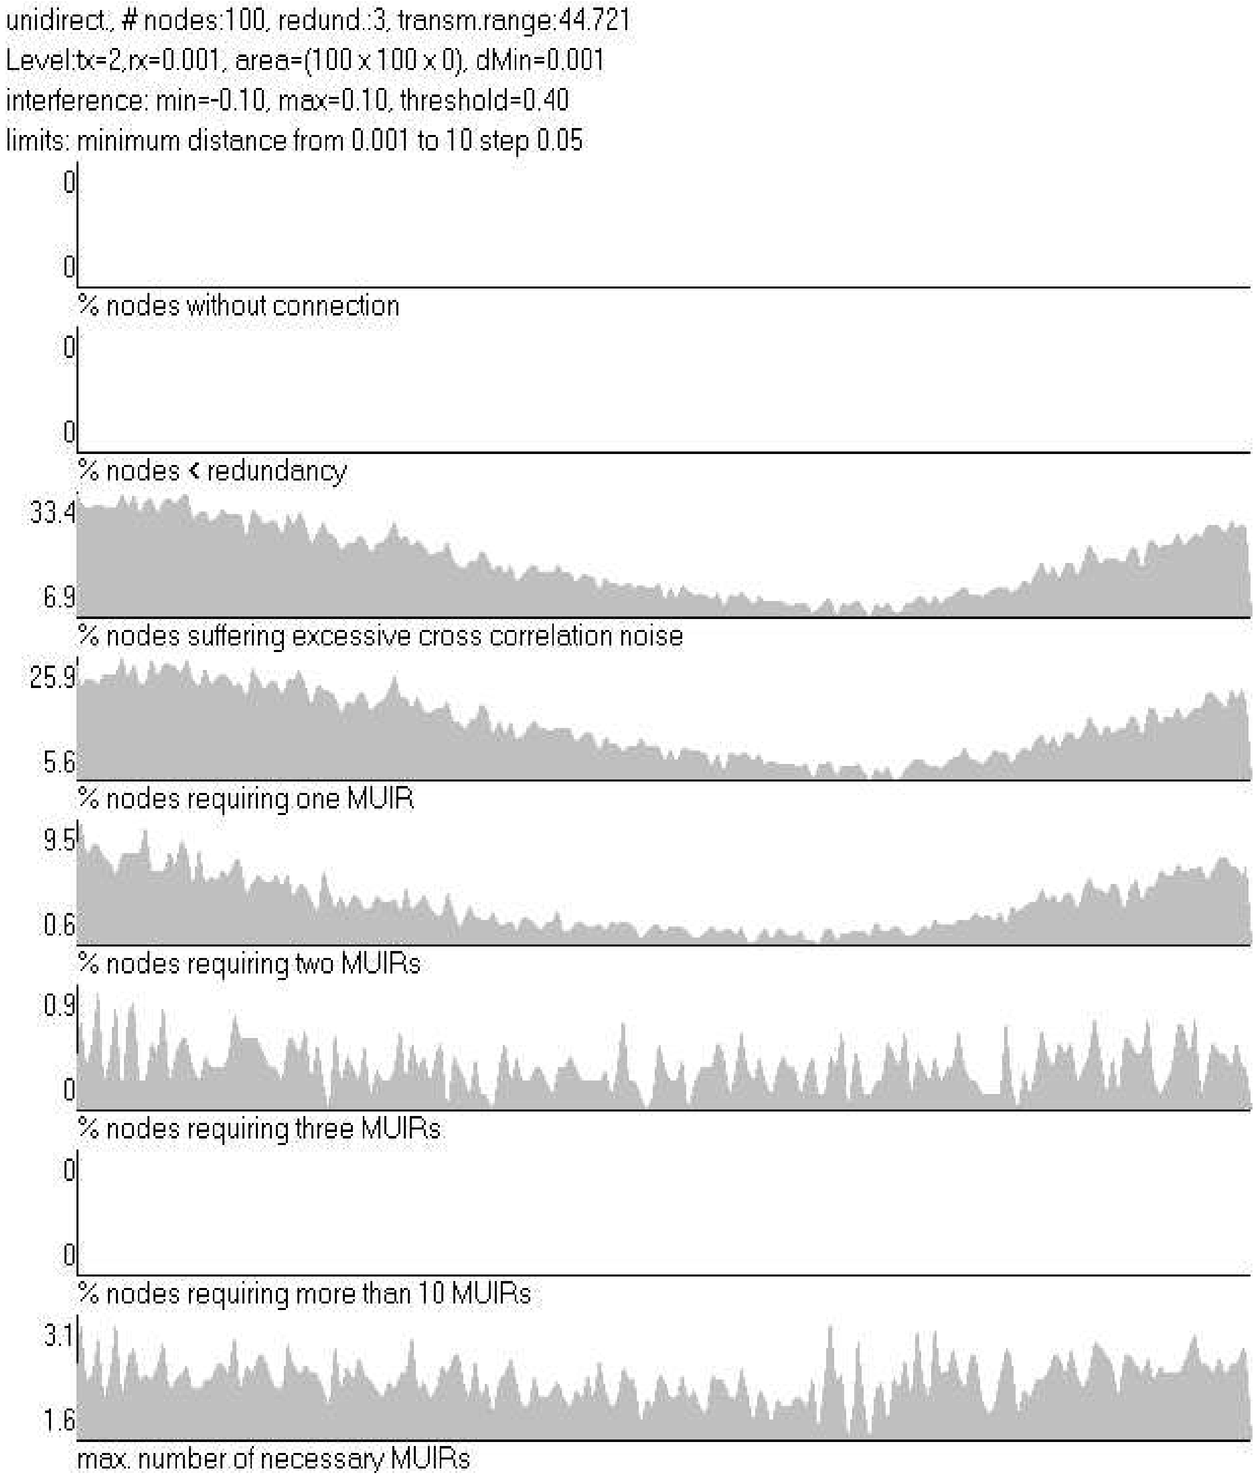

Figure 3: *Influence of Minimum Distance*

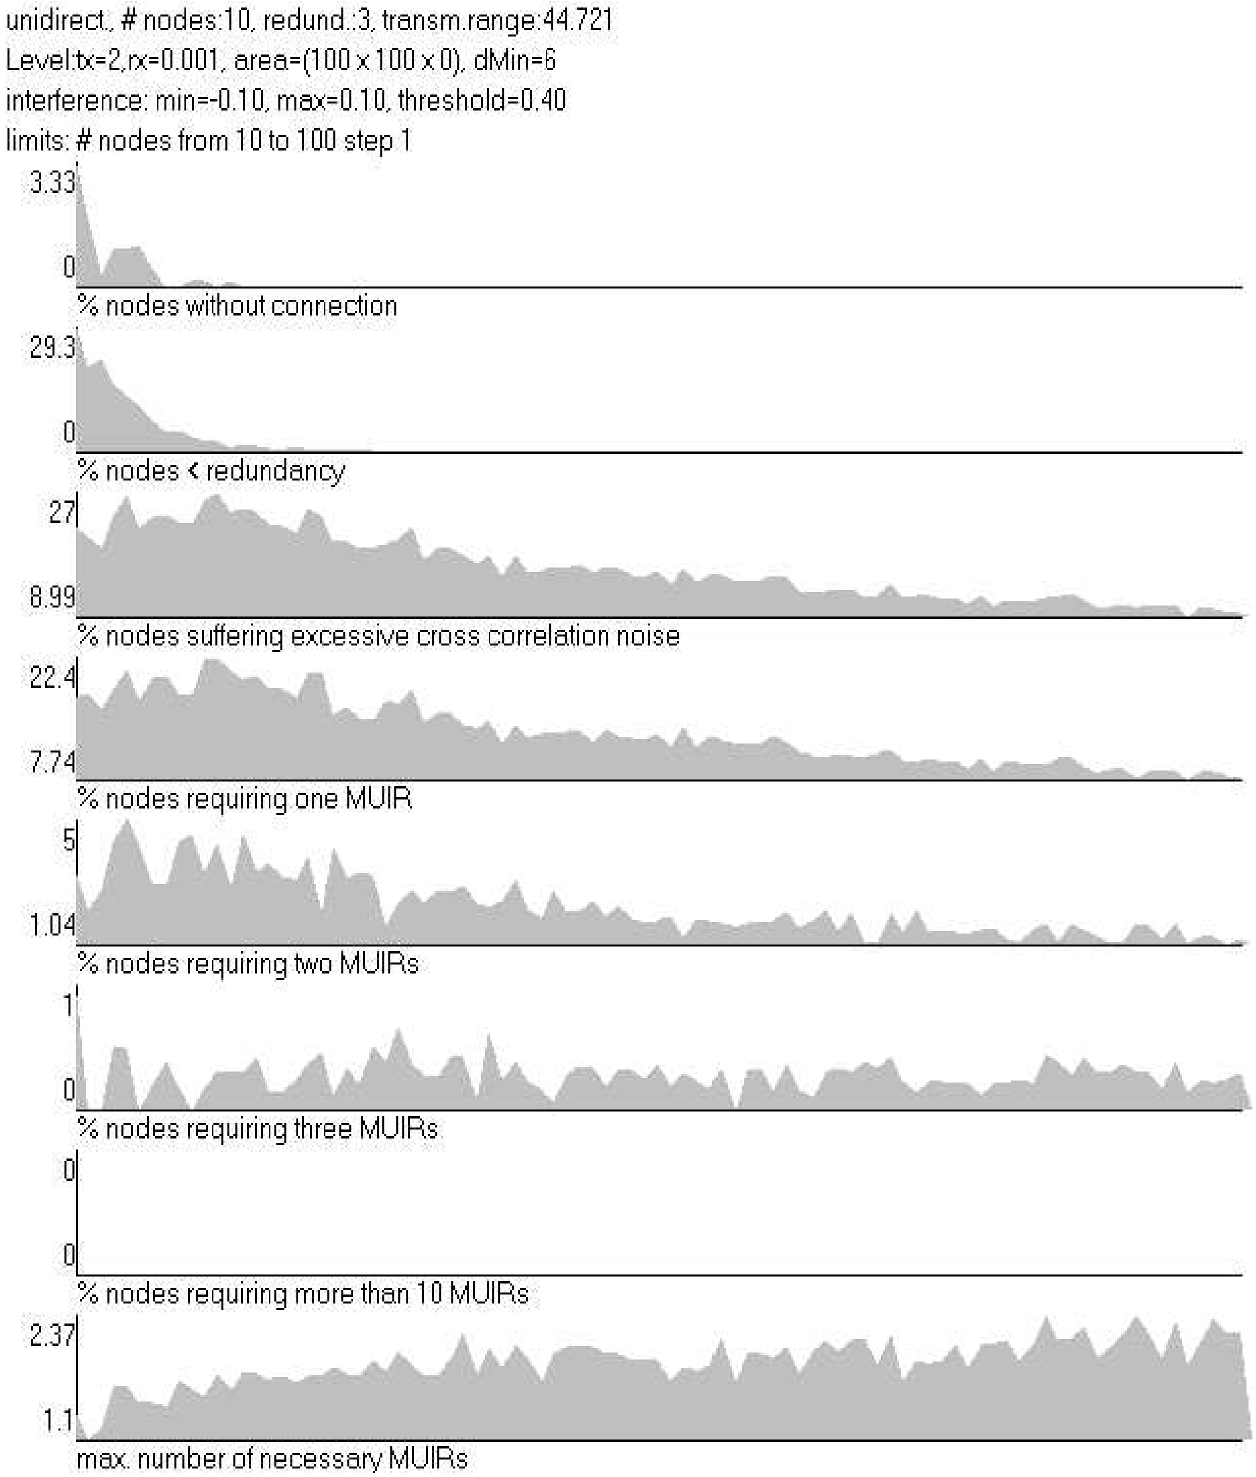

Figure 4: *Influence of Node Count on Connectivity*

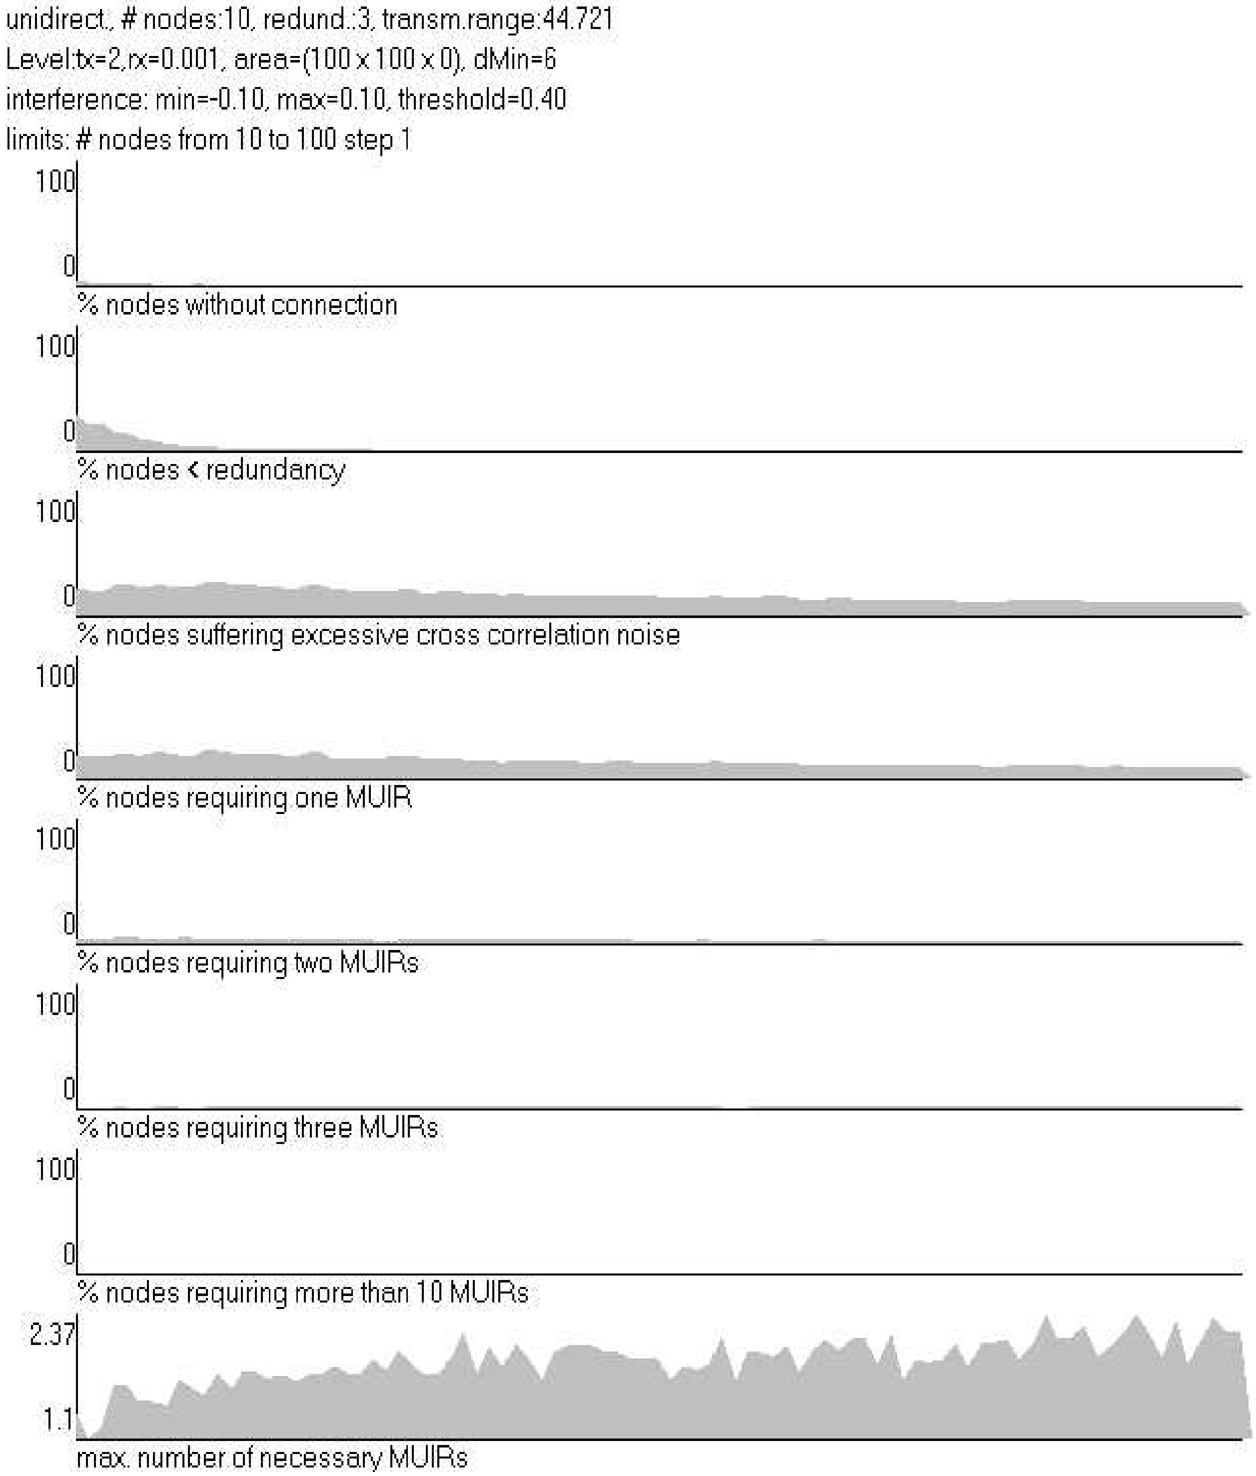

Figure 5: *Influence of Node Count on Connectivity (unscaled)*

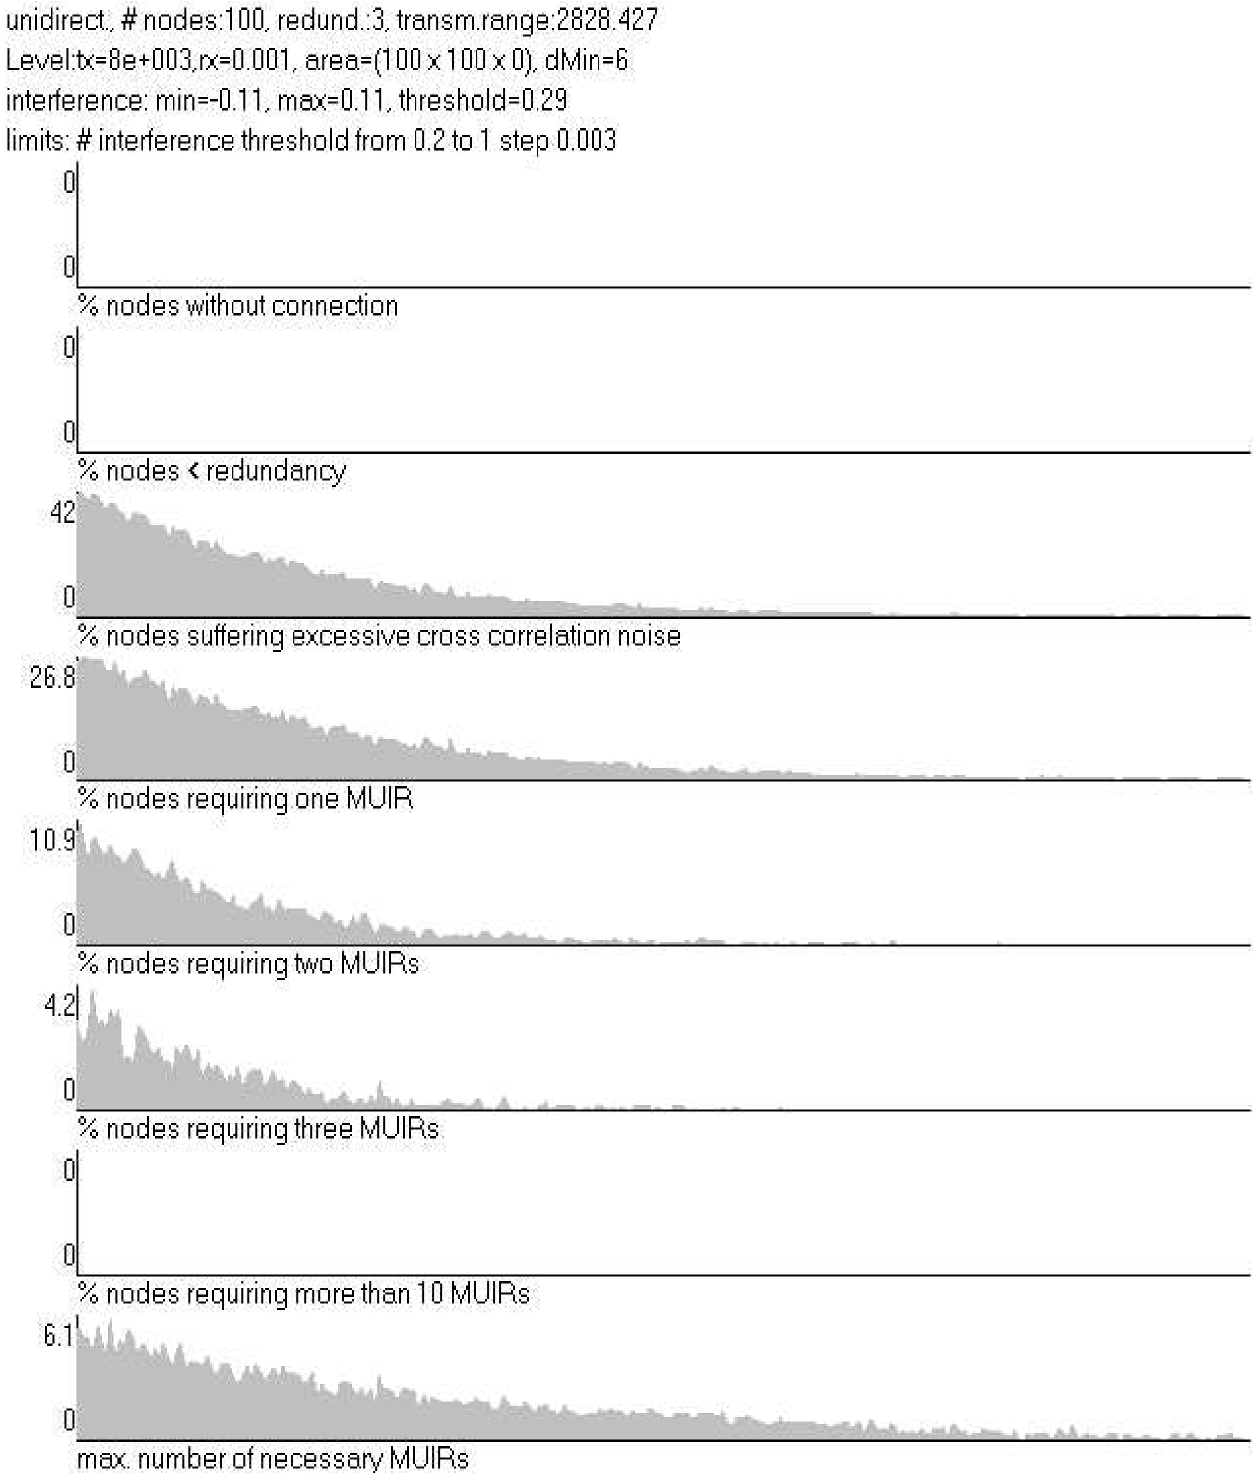

Figure 6: *Interference Threshold*

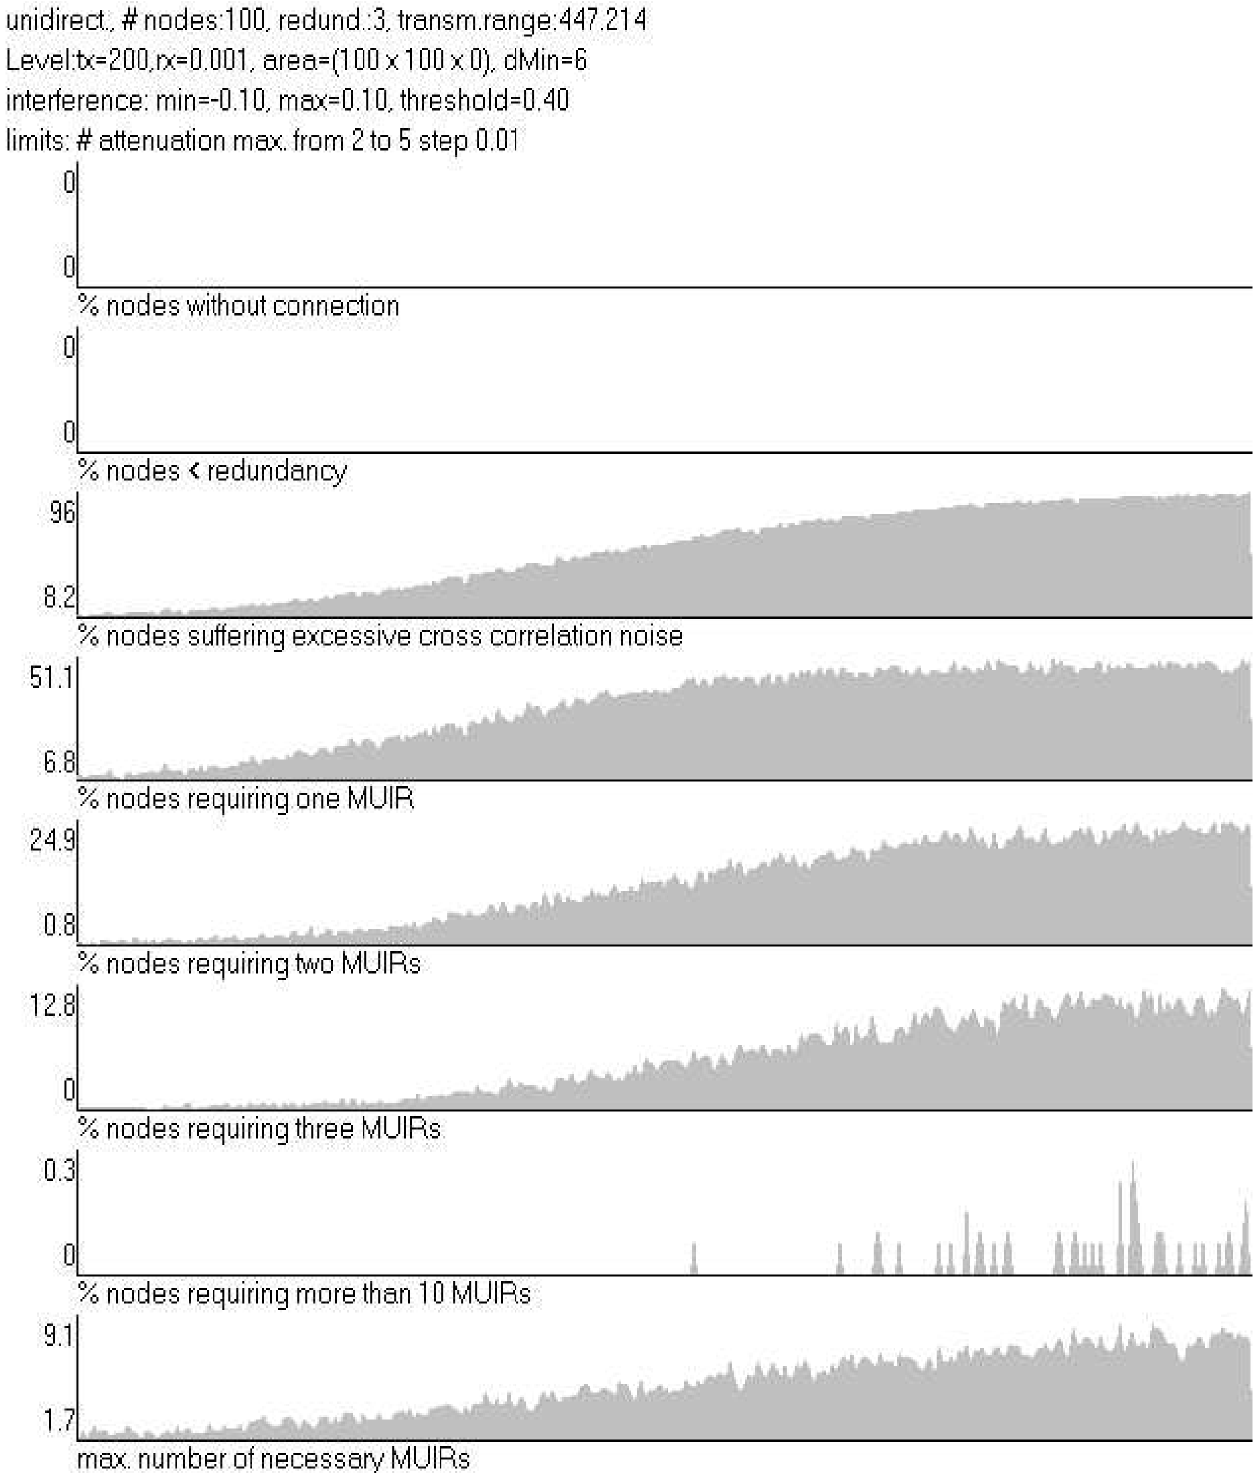

Figure 7: *Attenuation Maximum*

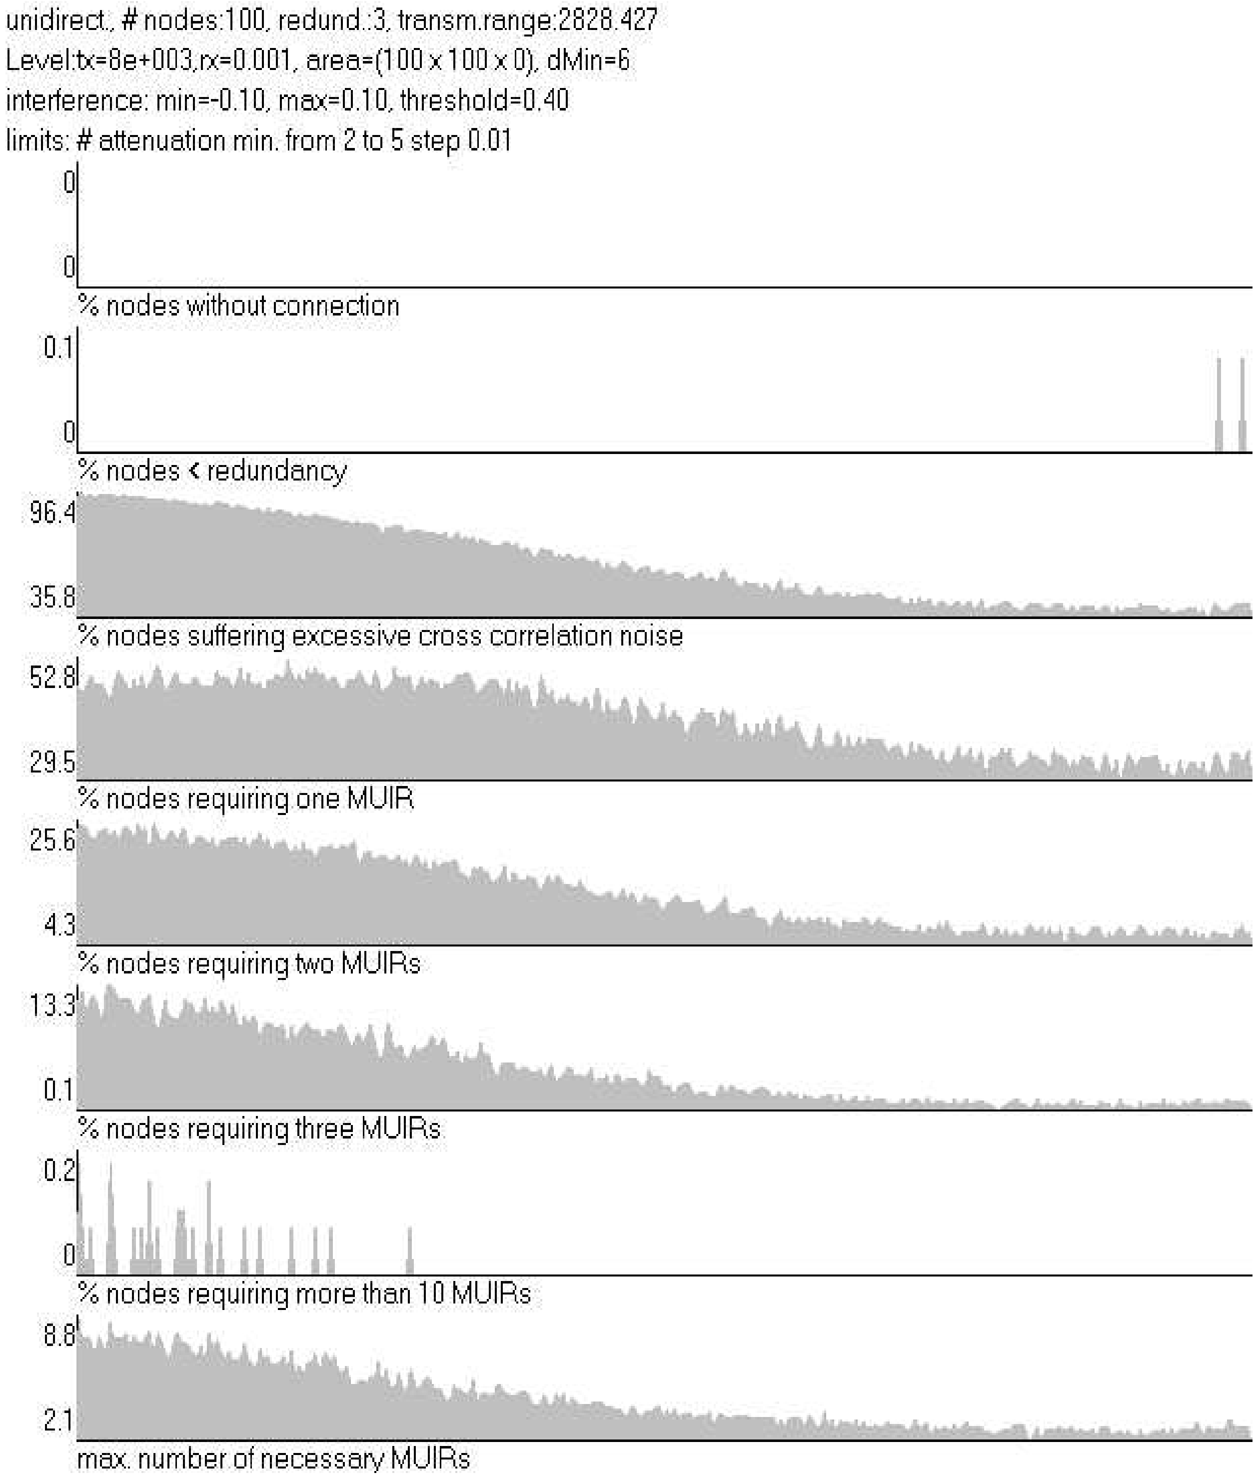

Figure 8: *Attenuation Minimum*

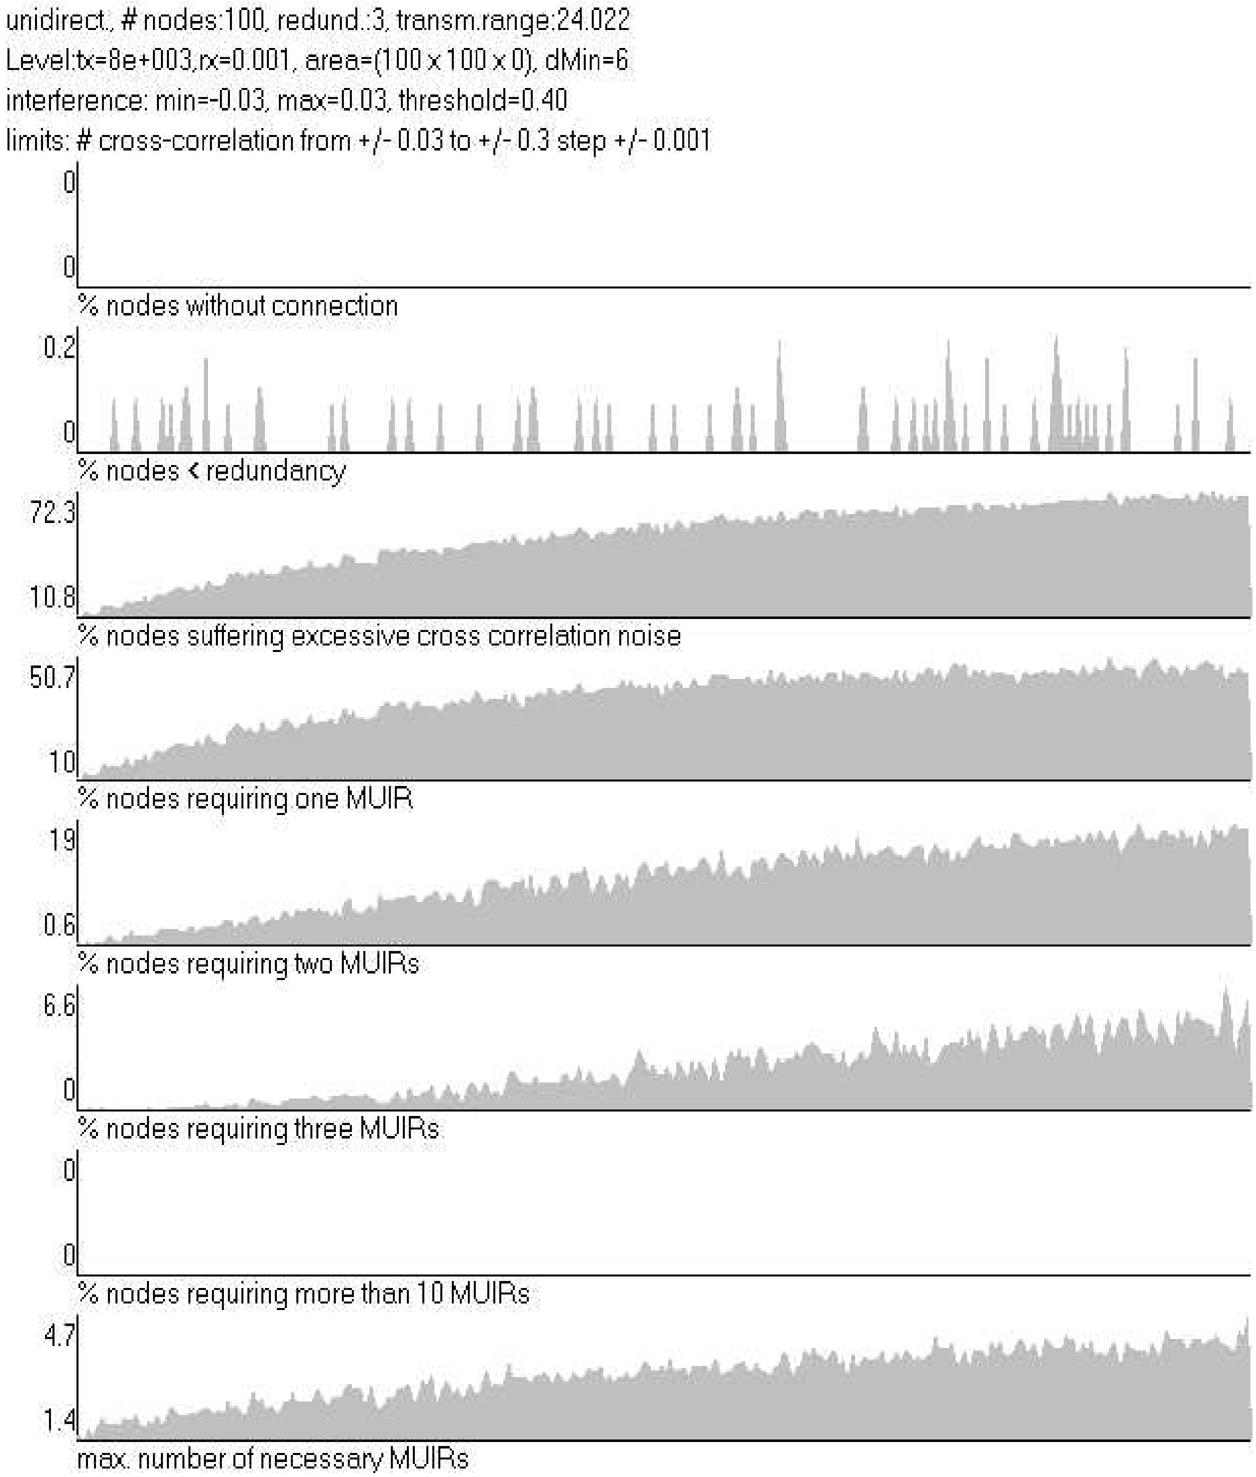

Figure 9: *Cross Correlation*