

Opportunistic Routing in Ad Hoc Networks: How many relays should there be? What rate should nodes use?

Joseph Blomer, Nihar Jindal
Dept. ECE, University of Minnesota, Minneapolis, MN, USA 55455,
Email: {blome015, nihar}@umn.edu

Abstract—Opportunistic routing is a multi-hop routing scheme which allows for selection of the best immediately available relay. In blind opportunistic routing protocols, where transmitters blindly broadcast without knowledge of the surrounding nodes, two fundamental design parameters are the node transmission probability and the transmission spectral efficiency. In this paper these parameters are selected to maximize end-to-end performance, characterized by the product of transmitter density, hop distance and rate. Due to the intractability of the problem as stated, an approximation function is examined which proves reasonably accurate. Our results show how the above design parameters should be selected based on inherent system parameters such as the path loss exponent and the noise level.

I. INTRODUCTION

Ad hoc networks operate on the basis of multi-hop routing, which allows information to be communicated across the network over a series of hops. End-to-end performance depends critically on the quality of the multi-hop routes used, but choosing good routes in a dynamic network (e.g. one where nodes are moving quickly) is a particularly difficult task. Opportunistic routing (e.g [1] - [6]) is well suited to such dynamic settings, as it can be performed with low overhead while also exploiting spatial diversity (in terms of fading and topology).

In this paper we focus on a purely opportunistic routing protocol, in which nodes broadcast packets in a completely blind manner (i.e. without any knowledge of surrounding nodes). More specifically, the protocol is described as follows:

- 1) In each time slot, each node randomly decides, with some fixed probability, to either transmit or receive.
- 2) A transmitting node will blindly broadcast a packet at some fixed spectral efficiency, whereas receivers listen for any transmission.
- 3) All receiver nodes that successfully decode the packet will send an ACK back to the transmitting node during an acknowledgement period. The ACK packets contain absolute geographic information about the RX location.
- 4) The transmitter selects the successful receiver which offers the most progress towards the packet's final destination and sends a message to that node electing it to become the next forwarder.

This protocol allows for the instantaneous choice of the best next hop, at the cost of an acknowledgement period and with the requirement for geographical information. Within this

algorithm, the key adjustable parameters are the *transmission probability* and the *transmission spectral efficiency*.

The transmission probability, denoted by p , determines the proportion of transmitters to receivers in each slot. When p is large, there is a large amount of interference between the nodes. This interference will cause only receivers close to transmitters to be potential forwarders. When the transmission probability is low, there are fewer simultaneous transmissions and thus less interference and more available relays. Thus longer hops are possible, but in fewer numbers. As a result, the trade-off with transmission probability is essentially between many simultaneous short hops, or fewer long hops.

When the spectral efficiency is high, a large signal-to-interference ratio (SIR) is required to decode. Thus, only relays which are close to the transmitter are likely to decode. On the other hand, a lower spectral efficiency allows nodes that are farther away to decode. Therefore, the trade-off with spectral efficiency is between shorter hops at higher data rate or longer hops at a lower data rate.

The objective of this paper is determining the transmission probability and spectral efficiency that optimally balance these trade-offs. We study blind opportunistic routing in a spatial model, as in [1]. In this model, end-to-end performance is characterized by the forward-rate-density, which is the product of average transmitter density, average hop distance and transmission rate (spectral efficiency). This is a function of measured network properties (e.g. path-loss exponent) and adjustable parameters (spectral efficiency and transmission probability). Our objective is to find the transmission probability and transmission spectral efficiency that maximize this metric. Studying this quantity directly is intractable, however we derive a reasonably accurate approximation and reinforce the conclusions with Monte Carlo simulations.

II. PRELIMINARIES

A. Network Model

Consider an infinite set of transmit/receive nodes Φ distributed according to a homogeneous 2-D Poisson point process (PPP) with density $\lambda \text{ m}^{-2}$. We consider a slotted transmission scheme. In each slot a node elects to become a transmitter with probability p , independent across users and slots. The set of locations of the transmitters (TX1, TX2, ..) denoted Φ^t and the set of locations of the receivers Φ^r

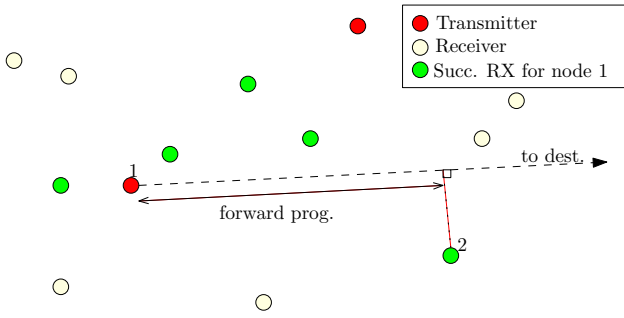


Fig. 1. Example network: TX node 1 has a message to be sent along dashed line. Among successful message receivers, RX node 2 offers the most forward progress and will forward the packet in the next slot.

then form independent PPP's of intensity λp and $\lambda(1 - p)$ respectively.

B. Channel Model

We consider a path-loss model with exponent $\alpha > 2$ and Rayleigh fading coefficients $h_{i,j}$ from TX i to RX j . Denoting the signal transmitted by TX i (located at X_i) as u_i , the received signal of RX j (at location X_j) is given by:

$$Y_j = \sum_i h_{i,j} |X_i - X_j|^{-\alpha/2} u_i + z. \quad (1)$$

For the time being, consider the network is interference limited, so thermal noise (z) is negligible. The case of a channel with thermal noise is considered in Sec IV.

Assuming the transmit powers are all equal, the SIR (signal to interference ratio) from TX i to RX j is defined:

$$S_{i,j} \triangleq \frac{|h_{i,j}|^2 |X_i - X_j|^{-\alpha}}{\sum_{k \neq i} |h_{k,j}|^2 |X_k - X_j|^{-\alpha}}. \quad (2)$$

We assume that all users transmit at rate equal to $R(\beta) = \log_2(1 + \beta)$; thus a communication is successful if and only if the received SIR is larger than a threshold value β .

III. PROGRESS-RATE PRODUCT

Although we are interested in the network-wide performance of the protocol, by the homogeneity of the PPP this is statistically equivalent to the performance of any transmitter. Thus, without loss of generality we can focus on TX0 located at the origin. From the perspective of TX0, all the other TX nodes form a homogeneous PPP with intensity λp (by Slivnyaks Theorem [7], the distribution is unaffected by conditioning on the presence of TX0). The transmitters and receivers are ordered by their distance to the origin i.e. TX1 is closest to the origin, TX2 the next, etc. For the sake of simplicity, the SIR from TX0 to RX j is denoted S_j and is given by $S_j = S_{0,j}$.

A. Forward progress density

We assume that the transmitter located at the origin has a packet to be sent to a receiver a very large distance away (e.g., fig. 1). Specifically, we assume the final destination is located an infinite distance away along the x -axis. Thus, the relay offering the most progress is the one with the maximum x -coordinate amongst the set of receivers that successfully decode TX0's packet. The corresponding progress, for a given network realization, is:

$$D(\Phi^r, \Phi^t, \beta) \triangleq \max_{X_j \in \Phi^r} \left[\mathbf{1}(S_j \geq \beta | \Phi^t) \left(|X_j| \cdot \cos(\theta(X_j)) \right) \right] \quad (3)$$

where $\mathbf{1}(S_j \geq \beta | \Phi^t)$ is the indicator function that the SIR at RX j (at point X_j) relative to TX0 (at the origin) is greater than the threshold. The expectation of this function with respect to the node process and fading is $d(p, \beta)$.

$$d(p, \beta) \triangleq \mathbb{E}^{\Phi, |h|^2} [D(\Phi^r, \Phi^t, \beta)] \quad (4)$$

This quantity is the expected progress (towards the destination) made in each hop. The forward progress density $f(p, \beta)$ then is:

$$f(p, \beta) \triangleq d(p, \beta) \lambda p. \quad (5)$$

Finally, the forward progress-rate-density function $K(p, \beta)$ is

$$\begin{aligned} K(p, \beta) &\triangleq (\lambda p) \cdot \mathbb{E}[D(\Phi^r, \Phi^t, \beta)] \cdot \log_2(1 + \beta) \\ &= f(p, \beta) \cdot R(\beta). \end{aligned} \quad (6)$$

Assuming each node in the network wishes to send data to another node a distance L away, at most $p \cdot d(p, \beta)/L$ packets can originate at each node in each slot [1]. Translating to bits/sec/Hz, this means that the maximum end-to-end data rate (per node) is

$$p \cdot d(p, \beta) \cdot \log_2(1 + \beta)/L \text{ [bps/Hz]}. \quad (7)$$

Thus, $K(p, \beta)$ is directly proportional to end-to-end rate.

The quantity $K(p, \beta)$, which is deterministic, is a function of transmission probability p and SIR threshold β . We seek the optimal values of β and p which maximize this product:

$$(p^*, \beta^*) = \arg \max_{p, \beta} K(p, \beta). \quad (8)$$

A closed-form expression for $K(p, \beta)$ cannot be found in general, so we rely on an approximation which is developed in the following section.

B. SIR Cell-Based Approximation

The basis of our approximation to $K(p, \beta)$ is the concept of an SIR cell [8], [9]. In the absence of fading, any RX within a particular region (i.e., the SIR cell) around a TX can decode, and the best forwarder is the RX in this region with the largest x -coordinate. Therefore, the forward progress is determined by the SIR cell – which is completely determined by the interferer locations and thus is a function only of the interferer process – and by the receiver locations (i.e., Φ^r). To reach a tractable approximation for $K(p, \beta)$, we remove the randomness in the

SIR cell and assume it is a deterministic region that is a function of p and β , while retaining the randomness in Φ^r .

We consider the average size of the SIR cell. This quantity is clearly defined in the absence of fading. On the other hand, the SIR cell is not well defined with fading because the SIR depends on independent fading RV's (specific to each receiver location) in addition to the interferer locations. However, we can derive a quantity analogous to the SIR cell area, denoted by v_0 , by integrating the point-wise success probability over all space. From [9], the success probability for a RX a distance y from the TX is:

$$\begin{aligned} p_0(y) &= \mathbb{P}^{\Phi^t}[S_j \geq \beta] \\ &= \exp\left(\frac{-\pi\lambda p|y|^2\beta^{\frac{2}{\alpha}}}{G(\alpha)}\right). \end{aligned} \quad (9)$$

Thus, v_0 is given by:

$$v_0 = \int_{y \in \mathbb{R}^2} p_0(y) dy = \frac{G(\alpha)}{\lambda p \beta^{\frac{2}{\alpha}}}. \quad (10)$$

where

$$G(\alpha) = \frac{\alpha}{2\Gamma(\frac{2}{\alpha})\Gamma(1 - \frac{2}{\alpha})} \quad (11)$$

and $\Gamma(z)$ is the gamma function $\Gamma(z) = \int_0^\infty t^{z-1} e^{-t} dz$.

We will then assume the SIR cell has area equal to its mean. We now approximate the SIR cell by a square centered on TX0 with side $\sqrt{v_0}$. In other words, we assume that all RX within this square are able to decode (ref. fig. 2). V_0^+ is the positive half of that square with area $v_0/2$. Under these assumptions, the expected forward progress is:

$$\begin{aligned} \tilde{d}_{sq}(p, \beta) &= \mathbb{E} \left[\max_{X_j \in V_0^+} x_j \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\max_{X_j \in V_0^+} x_j \mid \#(\{X_j \in V_0^+\}) = j \right] \right] \\ &= \sum_{j=0}^{\infty} \mathbb{E} \left[\max_{X_i \in V_0^+} x_i \mid \#(\{X_i \in V_0^+\}) = j \right] \times \\ &\quad \left(e^{-\lambda \cdot (1-p) \cdot \frac{v_0}{2}} \frac{(\lambda \cdot (1-p) \cdot \frac{v_0}{2})^j}{j!} \right) \\ &= \sum_{j=0}^{\infty} \left(\frac{\sqrt{v_0}}{2} \frac{j}{j+1} \cdot e^{-\lambda \cdot (1-p) \cdot \frac{v_0}{2}} \frac{(\lambda \cdot (1-p) \cdot \frac{v_0}{2})^j}{j!} \right) \\ &= \frac{\sqrt{v_0}}{2} \left(\frac{e^{-c} - (1-c)}{c} \right), \end{aligned} \quad (12)$$

where $\#(\cdot)$ gives the cardinality of a set and

$$c = \frac{\lambda \cdot (1-p) \cdot v_0}{2} = \frac{(1-p)}{p} \frac{G(\alpha)}{2\beta^{2/\alpha}}. \quad (13)$$

Note that the value c is the expected number of receivers in the half-square V_0^+ ; it is also equal to half the average number of successful outgoing transmissions from a TX as in [4].




Fig. 2. SIR cell is assumed to be a square V_0 of side length $\sqrt{v_0}$. The probability that the forward progress is x is the probability there are no RX in the dark shaded area.

Thus, our approximation to the progress-rate-density is:

$$f(p, \beta) \approx \tilde{f}_{sq}(p, \beta) = \tilde{d}_{sq}(p, \beta)(\lambda p) \quad (14)$$

$$K(p, \beta) \approx \tilde{K}_{sq}(p, \beta) = \tilde{f}_{sq}(p, \beta) \log_2(1 + \beta). \quad (15)$$

The optimal values then are the solution to the equation

$$\frac{\partial}{\partial p} \frac{\partial}{\partial \beta} \tilde{K}_{sq}(p, \beta) = 0. \quad (16)$$

Although there is no closed form solution to this equation, the approximation nonetheless yields valuable insight. By rearranging the terms in (12), we can interpret the forward progress approximation as the product of the maximum possible relay distance (i.e., half the side-length of the SIR cell) and a fractional term that captures how much of the maximum is attained by the best receiver.

$$\tilde{d}_{sq}(p, \beta) = \underbrace{\frac{\sqrt{v_0}}{2}}_{\text{max. relay dist.}} \underbrace{\left(1 - \frac{1 - e^{-c}}{c}\right)}_{\text{frac. of max.}} \quad (17)$$

The value v_0 is decreasing in p : as p increases, the amount of interference between transmitters is increasing, giving each transmitter a smaller expected SIR cell size. The value of c is also decreasing in p as a result of decreased v_0 and smaller receiver density.

The value $\tilde{K}_{sq}(p, \beta)$ can be written:

$$\begin{aligned} \tilde{K}_{sq}(p, \beta) &= \frac{1}{2} \log_2(1 + \beta) \sqrt{\lambda p G(\alpha) \beta^{-\frac{2}{\alpha}}} \left(1 - \frac{1 - e^{-c}}{c}\right) \end{aligned} \quad (18)$$

which lends the intuition that the end-to-end performance is increasing with p until the $1 - \frac{1 - e^{-c}}{c}$ fractional term becomes too small. When the value of c (expected potential forwarders) is greater than 3, this fractional term is greater than $2/3$. When p is large enough to cause c to go below 3, the performance drops off sharply.

IV. NUMERICAL RESULTS

For numerical results, Monte-Carlo simulations were performed. Locations of receivers and transmitters were realized using the properties of the PPP (c.f. [10]). SIR was calculated at each receiver and the largest progress offered was averaged.

A. Accuracy, Path-loss and Fading

Fig. 3 shows the accuracy of the approximation of the spatial density of progress (5), (14) for p and β . While the approximation is not absolutely accurate, it does correctly capture the dependence on p and β .

Fig. 4 shows the best p or β maximizing the progress-rate-density (15) for a fixed β or p , respectively. Again, the approximation helps to answer the essential question of what parameter choices are optimal.¹

In Fig. 5 the jointly optimal (p, β) are plotted for different values of the path loss exponent. Note that the optimal value of p is essentially the same for $\alpha = 3$ and 4, but that the optimizing SIR threshold increases with the path loss exponent. This increase is similar to [11], where the optimizing SIR threshold is found for a slightly different, but related, problem.

B. Sensitivity to parameter variation

Fig. 6 shows the normalized approximation of Eqn. (15) for path loss $\alpha = 3$ defined by

$$\tilde{K}_{\text{sq, norm}}(p, \beta) \triangleq \frac{\tilde{K}_{\text{sq}}(p, \beta)}{\max_{(p, \beta)} \tilde{K}_{\text{sq}}(p, \beta)}. \quad (19)$$

This figure shows how the performance decreases with different values of p and β . Observe that there exist a wide range of (p, β) pairs that are near optimal (within 90%). If there exists a strong reason to choose a particular β (e.g., limited code rates & modulations), or a particular p (e.g., power cycling, energy saving from different TX and RX powers), end-to-end performance does not suffer as long as the other parameter is appropriately chosen (to figs. 4, 6)

C. Robustness to Noise

In order to simulate the effect of noise on the forward progress, we modify the SIR (2) to include noise power σ^2 .

$$S_{i,j} = \frac{|h_{i,j}|^2 |X_i - X_j|^{-\alpha}}{\sum_{k \neq i} |h_{k,j}|^2 |X_k - X_j|^{-\alpha} + \sigma^2} \quad (20)$$

The metric for comparison is the average SNR to a node located at the average nearest-neighbor distance (d_{NN}).

$$\overline{\text{SNR}}_{NN} = \frac{\mathbb{E} [d_{NN}^2]^{-\frac{\alpha}{2}}}{\sigma^2} = \frac{(\lambda\pi)^{\alpha/2}}{\sigma^2} \quad (21)$$

Fig. 7 shows the maximum forward-rate-density is decreasing with increasing channel noise. Also, the number of potential forwarders is decreasing with increased channel noise.

Fig. 8 shows the effects of channel noise power on optimal p and β . The value of p^* is increasing and the value of β^* is decreasing with higher noise floor. The maximum possible spectral efficiency is expected to drop with increasing noise power. As explored in IV-B, when β is fixed to some value (in this case as a result of the noise floor), the transmission probability can increase to compensate.

¹Simulation results are similar for a non-fading environment.




Fig. 3. Spatial density of progress: approximation (thick curves) and simulation for $\lambda = 1$, $\alpha = 3$, varying SIR threshold β .

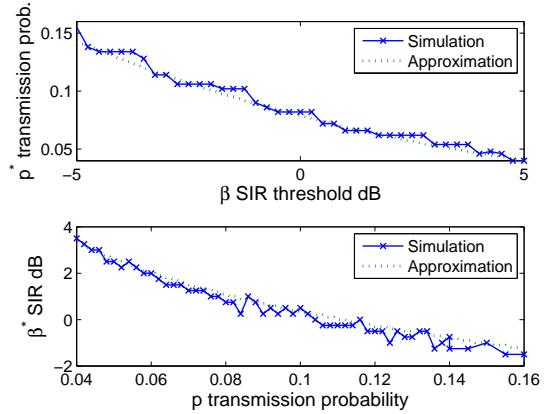


Fig. 4. Optimal p values giving best progress-rate-density for fixed values of SIR β (top). Optimal β values giving maximum progress-rate for fixed values p (bottom).

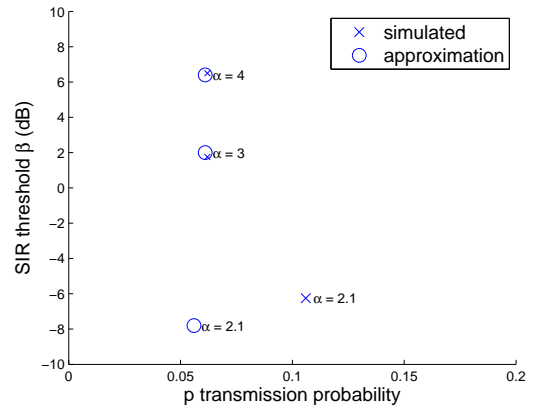


Fig. 5. Optimal choice of p and β giving maximum progress-rate for fixed α path-loss values.

V. CONCLUSION

We studied the case of opportunistic routing in an ad hoc wireless network with the goal of maximizing the product of forward progress density and the rate of data transmission. We developed an approximation using the concept of SIR cells and found for $\alpha = 3$ the optimal spectral efficiency is near 1.3 bps/Hz and the optimal probability of transmission is near 0.06. The results give the intuition that having a low transmission probability ensures a low amount of interference and a relatively high SIR threshold ensures only local transmissions are received.

REFERENCES

- [1] F. Baccelli, B. Blaszczyszyn, P. Muhlethaler, *An Aloha Protocol for Multihop Mobile Wireless Networks* IEEE Trans. on Inf. Th., Vol. 52, No. 2 Feb. 2006.
- [2] O.B.S. Ali, C. Cardinal, and F. Gagnon, *A performance analysis of multihop ad hoc networks with adaptive antenna array systems*.
- [3] S. Biswas R. Morris, *ExOR: Opportunistic Multi-Hop Routing for Wireless Networks* SIGCOMM Aug. 2005.
- [4] S. Weber, N. Jindal, R.K. Ganti, M. Haenggi *Longest edge routing on the spatial Aloha graph* IEEE Globecom 2008.
- [5] E.S. Sousa, J.A. Silvester *Optimum Transmission Ranges in a Direct-Sequence Spread-Spectrum Multihop Packet Radio Network* IEEE Jour. on Sel. Areas Comm Vol 8. No. 5 June 1990
- [6] M. Zorzi, S. Pupolin, *Optimum Transmission Ranges in Multihop Packet Radio Networks in the Presence of Fading* IEEE Trans of Comm. Vol. 43 No. 7 July 1995
- [7] D. Stoyan, W. Kendall, and J. Mecke, *Stochastic Geometry and Its Applications*. 2nd ed. New York: Wiley, 1996.
- [8] F. Baccelli, B. Blaszczyszyn, *Stochastic Geometry and Wireless Networks, Volume I: Theory* Foundations and Trends in Networking, NOW publishing, Vol 3, Iss. 3-4
- [9] F. Baccelli, B. Blaszczyszyn, *Stochastic Geometry and Wireless Networks, Volume II: Applications* Foundations and Trends in Networking, NOW publishing, Vol 4, Iss. 1-2
- [10] Martin Haenggi, *On Distances in Uniformly Random Networks*. IEEE Trans. on Inf. Th., Vol. 51, No. 10 Oct. 2005.
- [11] N. Jindal, J.G. Andrews, S. Weber, *Bandwidth Partitioning in Decentralized Wireless Networks* IEEE Trans. on Wireless Comm. Vol. 7 No. 12 Dec. 2008

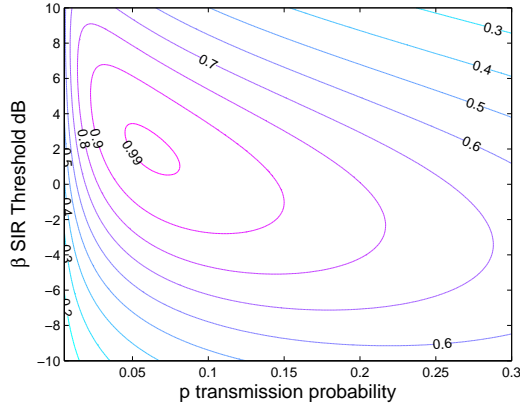


Fig. 6. Contour plot of normalized forward-rate-density function approximation.

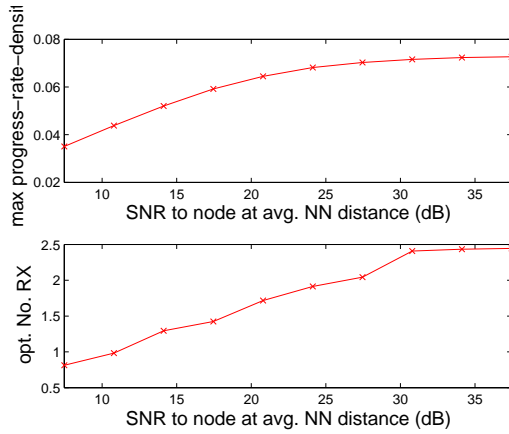


Fig. 7. Simulated best progress-rate-density achievable for varying SNR to node at mean nearest-neighbor distance $(\lambda\pi)^{\alpha/2}/\sigma^2$

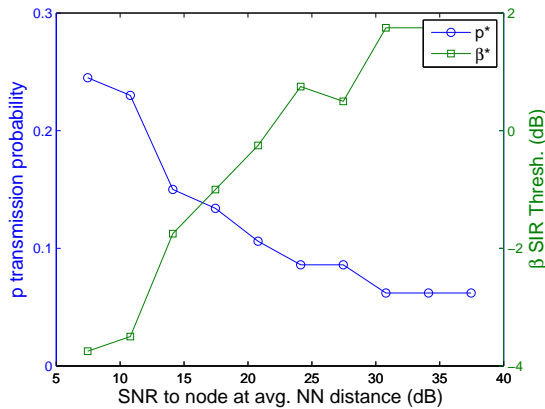


Fig. 8. Optimal values of p, β for varying SNR to node at mean nearest-neighbor distance $(\lambda\pi)^{\alpha/2}/\sigma^2$