The effectiveness of the linear hull effect

Sean Murphy

Communicated by Ed Dawson

Abstract. There is no linear hull effect in linear cryptanalysis.

1 Introduction

Linear cryptanalysis [9] is one of the standard techniques of assessing the security of block ciphers and is based on linear approximations to the plaintext, ciphertext and key. In our discussion of linear approximations, we consider an iterated block cipher encryption with plaintext p, ciphertext c and extended key k, where p, c and k are binary column vectors. We define the extended key k as the concatenation of the round subkeys derived from the block cipher key through the key schedule. A *linear approximation* in its most basic form is usually regarded as a statement of the form

$$\alpha^T \begin{pmatrix} p \\ c \end{pmatrix} = \alpha_P^T p + \alpha_C^T c = \gamma^T k \quad \text{with probability} \quad \frac{1}{2}(1 + \epsilon_\gamma).$$

The vectors α, α_P and α_C are known respectively as the (overall) data mask, the plaintext mask and the ciphertext mask, with $\alpha^T = (\alpha_P^T \; \alpha_C^T)$, and the vector γ as the key mask. The value ϵ_γ is known as the *imbalance* or *correlation* of the linear approximation, and $\frac{1}{2}\epsilon_\gamma$ as the *bias* of the linear approximation. If imbalance of the linear expression is nonzero, that is $\epsilon_\gamma \neq 0$, then the linear expression can potentially be used to give an estimate of one bit $\gamma^T k$ of key information. The number of plaintext–ciphertext pairs required to estimate this key bit to a required accuracy is proportional to ϵ_γ^{-2}. This is the procedure given by Algorithm 1 of [9].

The more sophisticated Algorithm 2 of [9] uses trial encryptions and decryptions of the outer rounds under various partial subkeys. The assumption of [9,11] is that the technique described by Algorithm 2 is equivalent to constructing a method

of distinguishing the distribution of the linear expression

$$\alpha^T \begin{pmatrix} p \\ c \end{pmatrix} + \gamma^T k,$$

which is zero with stated probability $\frac{1}{2}(1 + \epsilon_\gamma)$, from a uniform distribution. The number of plaintext–ciphertext pairs required to make this distinction to a required accuracy is then assumed to be proportional to ϵ_γ^{-2} (see [9]).

A widespread method of calculating the probability that such a single linear expression holds is to use the so-called *Piling-Up Lemma* of [9]. In the analysis of many block ciphers, such use of the Piling-Up Lemma can generate a number of linear expressions with the same data mask but differing key masks, that is a key mask set Γ and a collection of expressions of the form

$$\alpha^T \begin{pmatrix} p \\ c \end{pmatrix} = \gamma^T k \quad \text{with probability} \quad \frac{1}{2}(1 + \epsilon_\gamma) \quad [\gamma \in \Gamma].$$

However, we show in Section 3 that a collection of expressions of the above form is not in general well-defined. It is simply not sustainable to use the Piling-Up Lemma to make meaningful probability calculations for such a collection of expressions.

The supposed existence of a number of such expressions with $\epsilon_\gamma \neq 0$ generated by the Piling-Up Lemma appears nevertheless to be the motivation for the concept of the linear hull of linear expressions introduced by [11]. The linear hull for data mask α is considered by [11] to be the set of all such above expressions for different key masks γ. It is asserted by [11] that existence of such a linear hull containing many such unbiased expressions generally increases the efficiency of Algorithm 2 of [9]. Furthermore, such an assertion appears to be generally accepted and widely used in the analysis of block ciphers. For example, a standard reference work on cryptology makes the following statement about the use of the linear hull in cryptanalysis [3].

> *Linear hulls:* Estimating the bias of approximations by constructing linear characteristics is very convenient, but in some cases, the value derived in this way diverges significantly from the actual bias. The most important cause for this difference is the so-called *linear hull effect*, first described by Nyberg in 1994 [11]. The effect takes place when the correlation between plaintext and ciphertext bits, described by a specific linear approximation, can be explained by multiple linear characteristics, each with a non-negligible bias, and each involving a different set of key bits. Such a set of linear characteristics with identical input and output masks is called a *linear hull*. Depending on the

value of the key, the different characteristics will interfere constructively or destructively, or even cancel out completely. If the set of keys used in different linear characteristics are independent, then this effect might considerably reduce the average bias of [a single linear] expression, and thus the success rate of the simple attack described above [Algorithm 1 of [9]]. Nyberg's paper [11] shows, however, that the more efficient attacks (Algorithm 2) described in [9], which only use the linear approximations as a distinguisher, will typically benefit from the linear hull effect.

We show that this so-called linear hull effect in the sense of [3, 11] does not exist.

2   The fundamental probability

Our analysis of the linear hull effect is based on the fundamental probability (Definition 2.1), together with the related concept of the fundamental imbalance (Definition 2.2). The fundamental probability is a well-defined probability for use in linear cryptanalysis, particularly in analysing the linear hull effect.

Definition 2.1. The *fundamental probability* $q(k)$ of data mask α for block cipher encryptions under the fixed extended key k is

$$q(k) = \mathbf{P}_k \left(\alpha^T \begin{pmatrix} p \\ c \end{pmatrix} = 0 \right).$$

Definition 2.2. The *fundamental imbalance* of data mask α under fixed key k with fundamental probability $q(k) = \frac{1}{2}(1 + \eta(k))$ is

$$\eta(k) = \mathbf{P}_k \left(\alpha^T \begin{pmatrix} p \\ c \end{pmatrix} = 0 \right) - \mathbf{P}_k \left(\alpha^T \begin{pmatrix} p \\ c \end{pmatrix} = 1 \right) = 2q(k) - 1.$$

We note that the fundamental imbalance can be considered as a *correlation*, as the term is defined by [5, 6], parameterised by the key. Methods for calculating the fundamental imbalance without using the Piling-Up Lemma in certain situations are given by [5, 6] and by [12, Theorem 7].

3   Probabilistic interpretation of the linear hull

We noted that the fundamental probability (Definition 2.1) is well-defined. However, the probability statements made in the usual definition of a linear hull are not

in general well-defined. To demonstrate this, we consider very carefully the nature of the statements about linear expressions used in the linear hull for a data mask α.

In the standard formulation of linear cryptanalysis given in Section 1, a linear approximation is defined directly in terms of a key mask γ, that is, a linear approximation is a statement of the form

$$\alpha^T \begin{pmatrix} p \\ c \end{pmatrix} = \gamma^T k \quad \text{with probability} \quad \frac{1}{2}(1 + \epsilon_\gamma),$$

which only depends on the key through the value of $\gamma^T k$. In terms of the fundamental probability for data mask α, we have

$$\mathbf{P}_k \left(\alpha^T \begin{pmatrix} p \\ c \end{pmatrix} = 0 \right) = \frac{1}{2}(1 + \eta(k)) = \begin{cases} \frac{1}{2}(1 + \epsilon_\gamma) & [\gamma^T k = 0], \\ \frac{1}{2}(1 - \epsilon_\gamma) & [\gamma^T k = 1]. \end{cases}$$

Thus the fundamental imbalance $\eta(k)$ is given in terms of the usual form of the imbalance ϵ_γ used in linear cryptanalysis by $\eta(k) = (-1)^{\gamma^T k} \epsilon_\gamma$. However, for a fixed key k, the fundamental imbalance $\eta(k)$ is clearly a constant. Thus for the above probability to be well-defined, we require $(-1)^{\gamma^T k} \epsilon_\gamma$ to be constant for all non-trivial key masks γ for fixed k. We now discuss this issue in more detail.

The linear hull is usually considered to be defined by several probabilistic statements such as (for $\gamma \neq \gamma'$)

$$\alpha^T \begin{pmatrix} p \\ c \end{pmatrix} = \gamma^T k \quad \text{with probability} \quad \frac{1}{2}(1 + \epsilon_\gamma)$$

and

$$\alpha^T \begin{pmatrix} p \\ c \end{pmatrix} = \gamma'^T k \quad \text{with probability} \quad \frac{1}{2}(1 + \epsilon_{\gamma'}).$$

According to these linear expressions, for a fixed key k, the single bit $\alpha^T \begin{pmatrix} p \\ c \end{pmatrix}$ has to simultaneously take the value $\gamma^T k$ with probability $\frac{1}{2}(1 + \epsilon_\gamma)$ and the value $\gamma'^T k$ with probability $\frac{1}{2}(1 + \epsilon_{\gamma'})$. This means there are four cases required to evaluate the fundamental probability, as shown in Table 1.

We can therefore deduce that $\epsilon_\gamma = \epsilon_{\gamma'} = 0$. Thus the probability statements used in specifying a linear hull in the sense of [3, 11] are not well-defined if there is more than one non-trivial linear approximation in the linear hull. A linear hull in the original and widespread sense of [3, 11], that is a collection of several linear expressions each with a significant imbalance or bias, is not a well-defined probabilistic concept.

	$\gamma'^T k = 0$	$\gamma'^T k = 1$
$\gamma^T k = 0$	$\frac{1}{2}(1 + \epsilon_\gamma) = \frac{1}{2}(1 + \epsilon_{\gamma'})$	$\frac{1}{2}(1 + \epsilon_\gamma) = \frac{1}{2}(1 - \epsilon_{\gamma'})$
$\gamma^T k = 1$	$\frac{1}{2}(1 - \epsilon_\gamma) = \frac{1}{2}(1 + \epsilon_{\gamma'})$	$\frac{1}{2}(1 - \epsilon_\gamma) = \frac{1}{2}(1 - \epsilon_{\gamma'})$

Table 1. Fundamental probability $\mathbf{P}_k(\alpha^T (p^T \ c^T)^T = 0)$.

4 Analogues in differential cryptanalysis of linear hulls

A much-repeated heuristic justification for the linear hull effect is given in [11] by claiming that the use of the linear hull is analogous to the use of the *differential* [8] in differential cryptanalysis [1,2]. However, a detailed examination shows that this analogy is not sustainable. We therefore state the claim of [11] of an analogy between differential and linear cryptanalysis using our notation.

> We conclude that Algorithm 2 [of [9]] makes in fact use of a family of linear approximate expressions

$$\alpha_P p + \alpha_C^T c + \gamma^T k = \alpha^T \begin{pmatrix} p \\ c \end{pmatrix} + \gamma^T k$$

> where [the data mask] α is fixed but [the key mask] γ varies. This means that the round approximations which uniquely determine γ and are uniquely determined by γ, can be chosen in all possible ways to form a chain of approximations from $\alpha_P^T p$ to $\alpha_C^T c$. Hence there is a close analog with what is called differentials in differential cryptanalysis [8].

The key mask γ does indeed determine a series of round data masks with the outer round masks constrained by α. These data masks give a random process [4,7] with a state space consisting of two elements $\{0, 1\}$, in for example the manner described by [10]. The usual method of analysing this random process in linear cryptanalysis is to use the Piling-Up Lemma [9], which is applicable to this random process if it is a Markov process [10]. However, whether or not the Piling-Up Lemma is applicable, both states are considered by the Piling-Up calculation. Loosely speaking, the Piling-Up Lemma probability calculation may incorrectly "assign probability" to the states if the Markov assumption is not valid, but "all probability is assigned". There are no "ignored" states in linear cryptanalysis, and there is certainly no "missing probability" not considered by the Piling-Up Lemma waiting to be found.

In differential cryptanalysis, a pair of plaintexts is encrypted under the same fixed key. A random process is derived by taking the difference of those plaintexts

at every round, so giving a random process with a state space of size 2^n. Under the assumption that this differential cryptanalysis random process is a Markov process, the *differential* [8] is calculated by using the product of the matrices of one-step (round) transition probabilities. The *characteristic* [1,2] is essentially obtained by setting all but one element to 0 in each of these one-step transition matrices and then calculating the product of these revised matrices. A value for a probability calculated using a differential [8] is therefore always at least that of the same value for a probability calculated using a characteristic [1,2]. Thus there is a meaningful sense in differential cryptanalysis in which the differential can be thought of as "finding probabilities" missed by the characteristic, as the characteristic "ignores" most of the states in the random process.

The correct analogue of the linear hull in differential cryptanalysis is the *enhanced characteristic*, considered in [1, Sections 5.1, 6.5]. The enhanced characteristic gives an enhanced differential cryptanalysis random process in which the state at any round is the difference between the data values at that round (as for standard differential cryptanalysis) and also a collection of data values at certain bit positions. Similarly, in linear cryptanalysis, the simultaneous use of two key masks with the same data mask, really defines sequences of 2-dimensional data masks for the inner rounds. This gives rise to a random process with, in general, a 4-element state space rather than the 2-element state space given by a single key mask. For both the enhanced characteristic in differential cryptanalysis and the linear hull in linear cryptanalysis, the state space of the underlying random process state space is obtained by refining the standard differential cryptanalysis or linear cryptanalysis state space.

The linear hull therefore defines a new (enhanced) random process beyond the standard linear cryptanalysis random process. By contrast, probability calculations in differential cryptanalysis using either characteristics or differentials are conducted with respect to the same standard differential cryptanalysis random process. Thus a calculation using a linear hull and a calculation using a differential are fundamentally different in relation to their underlying standard random process. A differential can find "unused probability", whereas a linear hull simply cannot. There is no analogue between linear hulls and differentials in the manner claimed by [11].

5   Supposed data requirements for the linear hull effect

The linear hull effect supposedly reduces the number of plaintext–ciphertext pairs required to distinguish to a given accuracy a given distribution from the uniform distribution [3, 11]. Before analysing this claim, we first calculate the data re-

quirements exactly under the assumptions of [9, 11]. For a fixed extended key k, the number N_k of plaintext–ciphertext pairs required to use the distinguisher to a required accuracy is stated [9, 11] to be asymptotically inversely proportional to $(q(k) - \frac{1}{2})^2$. Under the assumption that all extended keys are equally likely, then the mean number (over all keys) N of plaintext–ciphertext pairs required to use the distinguisher to a specified degree of accuracy is given by $N = \mathbf{E}[N_k]$, for some suitably-defined expectation $\mathbf{E}$. If we let $\mathcal{K}$ denote the set of extended keys, then the mean number N of plaintext–ciphertext pairs required would then be proportional to

$$\mathbf{E}\left[\left(\left(q(k) - \tfrac{1}{2}\right)^2\right)^{-1}\right] = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \left(\tfrac{1}{4}\eta(k)^2\right)^{-1}.$$

We now consider the data requirement for the use of the linear hull given by [11]. The supposed existence of a linear hull effect depends on the assertion that this data requirement can be expressed in terms of a quantity defined by the *Fundamental Theorem* of [11]. We now examine this assertion. Accordingly, we let

$$r(\gamma, k) = \frac{1}{2}\left(1 + (-1)^{\gamma^T k}\eta(k)\right),$$

so $r(\gamma, k)$ is the quantity referred to as "$p(a, b, c; k)$" in [11], where "a, b" refer to the data mask α and "c" refers to the key mask γ. We note that "$p(a, b, c; k)$" is referred to as a probability by [11], but is not in general a well-defined probability (Section 3). The Fundamental Theorem considers the quantity ψ given by

$$\psi = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \left(r(\gamma, k) - \tfrac{1}{2}\right)^2 = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \frac{1}{4}\eta(k)^2 = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \frac{1}{4}\left(q(k) - \tfrac{1}{2}\right)^2.$$

Thus the quantity ψ considered by the Fundamental Theorem can be expressed in terms of the above expectation $\mathbf{E}$ as

$$\psi = \mathbf{E}\left[\left(q(k) - \tfrac{1}{2}\right)^2\right] = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \frac{1}{4}\eta(k)^2.$$

The claim for the reduction in data complexity given by the linear hull effect is then based on the assertion that the number N of plaintext–ciphertext pairs required to distinguish the distribution from uniform is proportional to ψ^{-1}, that is to say proportional to

$$\psi^{-1} = \mathbf{E}\left[\left(q(k) - \tfrac{1}{2}\right)^2\right]^{-1} = \left(\frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \frac{1}{4}\eta(k)^2\right)^{-1}.$$

We now compare these two quantities for distinguishing the specified distribution from a uniform distribution. The actual expected number N of plaintext–ciphertext pairs and the number given by the application of the Fundamental Theorem in the manner of [11] are respectively proportional to

$$\mathbf{E}\left[\left(\left(q(k)-\tfrac{1}{2}\right)^2\right)^{-1}\right] \quad \text{and} \quad \mathbf{E}\left[\left(q(k)-\tfrac{1}{2}\right)^2\right]^{-1}.$$

These two quantities can easily be compared by using Jensen's inequality [13], which we state in Lemma 5.1.

Lemma 5.1 (Jensen's inequality). *A random variable X and convex function ζ satisfy $\zeta(\mathbf{E}[X]) \leq \mathbf{E}[\zeta(X)]$.*

Inversion of the positive real numbers is a convex function, so Jensen's inequality gives

$$\mathbf{E}\left[\left(\left(q(k)-\tfrac{1}{2}\right)^2\right)^{-1}\right] \geq \mathbf{E}\left[\left(q(k)-\tfrac{1}{2}\right)^2\right]^{-1}.$$

The supposed data requirement under the assumptions of [9, 11] is proportional to the left-hand side of the above inequality, whereas the quantity ψ considered by the Fundamental Theorem only addresses the right-hand side of this inequality. Thus the Fundamental Theorem can only ever be used to give a lower bound on the data complexity even under the Algorithm 2 assumptions of [9, 11]. We illustrate this point in Examples 5.2 and 5.3.

Example 5.2. We consider a linear cryptanalysis for data mask α where the fundamental probability is given by

$$q(k) = \mathbf{P}_k\left(\alpha^T\begin{pmatrix} p \\ c \end{pmatrix} = 0\right) = \frac{1}{2}(1 + (-1)^{\gamma^T k}\epsilon_\gamma + (-1)^{\gamma'^T k}\epsilon_{\gamma'}),$$

where $\epsilon_\gamma, \epsilon_{\gamma'} \neq 0$. Thus the fundamental imbalance is given by

$$\eta(k) = (-1)^{\gamma^T k}\epsilon_\gamma + (-1)^{\gamma'^T k}\epsilon_{\gamma'}.$$

The data requirement under the assumptions of [9, 11] to distinguish this distribution from a uniform distribution to a given degree of accuracy is proportional to

$$\mathbf{E}\left[\left(\left(q(k)-\tfrac{1}{2}\right)^2\right)^{-1}\right] = \frac{1}{4}\left(\frac{2}{(\epsilon_\gamma + \epsilon_{\gamma'})^2} + \frac{2}{(\epsilon_\gamma - \epsilon_{\gamma'})^2}\right) = \frac{\epsilon_\gamma^2 + \epsilon_{\gamma'}^2}{(\epsilon_\gamma^2 - \epsilon_{\gamma'}^2)^2}.$$

By contrast, the approach of [11] based on the Fundamental Theorem asserts that this supposed data requirement is proportional to

$$\mathbf{E}\left[\left(q(k) - \frac{1}{2}\right)^2\right]^{-1} = \frac{1}{4}\left(2(\epsilon_\gamma + \epsilon_{\gamma'})^2 + 2(\epsilon_\gamma - \epsilon_{\gamma'})^2\right)^{-1} = \frac{1}{\epsilon_\gamma^2 + \epsilon_{\gamma'}^2}.$$

However,

$$(\epsilon_\gamma^2 + \epsilon_{\gamma'}^2)(\epsilon_\gamma^2 - \epsilon_{\gamma'}^2)^{-2} > (\epsilon_\gamma^2 + \epsilon_{\gamma'}^2)^{-1},$$

as $(\epsilon_\gamma^2 + \epsilon_{\gamma'}^2)^2 > (\epsilon_\gamma^2 - \epsilon_{\gamma'}^2)^2$, so we can see that

$$\mathbf{E}\left[\left(\left(q(k) - \frac{1}{2}\right)^2\right)^{-1}\right] > \mathbf{E}\left[\left(q(k) - \frac{1}{2}\right)^2\right]^{-1}.$$

For this example, we can easily see that the data requirement under the assumptions of [9, 11] averaged over all extended keys is always larger than that asserted by the supposed linear hull effect.

Example 5.3. We consider the cryptanalysis of Example 5.2 for a block cipher where $\epsilon_{\gamma'} = \epsilon_\gamma$. Such a block cipher would be considered to possess a large linear hull effect in the "linear hull literature". The fundamental probability $q(k)$ for data mask α is given by

$$q(k) = \begin{cases} \frac{1}{2}(1 + (\epsilon_\gamma + \epsilon_{\gamma'})) & [\gamma^T k = \gamma'^T k = 0], \\ \frac{1}{2}(1 - (\epsilon_\gamma + \epsilon_{\gamma'})) & [\gamma^T k = \gamma'^T k = 1], \\ \frac{1}{2} & [\gamma^T k \neq \gamma'^T k]. \end{cases}$$

Clearly, for any key $k \notin \ker(\gamma + \gamma')^T$, that is for half of all the extended keys, the distribution in question is uniform, so by definition is indistinguishable from a uniform distribution. This is reflected in that fact that the data requirement under the assumptions of [9, 11], proportional to $(\epsilon_\gamma^2 + \epsilon_{\gamma'}^2)(\epsilon_\gamma^2 - \epsilon_{\gamma'}^2)^{-2}$, is formally infinite. By contrast, the supposed linear hull effect gives this data requirement as being proportional to $(\epsilon_\gamma^2 + \epsilon_{\gamma'}^2)^{-1}$, a finite quantity.

6   Conclusions

The usual method of quantifying the supposed linear hull effect assumes that expectation and inversion are operations on a random variable which commute. Jensen's inequality shows that this assumption is generally incorrect. Thus the supposed linear hull effect simply ignores Jensen's inequality, a fundamental result in probability theory and statistical inference. Furthermore, Jensen's inequality

shows that the Fundamental Theorem of [11] can only ever give a lower bound for the data requirement under the assumptions of [9, 11] for using a collection of "linear approximations" with the same data mask.

The linear hull effect, in the usual sense [3, 11] of typically improving the average efficiency of Algorithm 2 of [9] under the assumptions given in this paper, is an illusion.

Bibliography

[1] E. Biham and A. Shamir, Differential cryptanalysis of the DES-like cryptosystems, *J. Cryptology* **4** (1991), 3–72.

[2] E. Biham and A. Shamir, *Differential Cryptanalysis of the Data Encryption Standard*, Springer, Berlin, 1993.

[3] A. Biryukov and C. De Cannière, Linear cryptanalysis for block ciphers, in: *Encyclopedia of Cryptography and Security*, Springer, Berlin (2005), 351–354.

[4] D. R. Cox and H. D. Miller, *The Theory of Stochastic Processes*, Chapman and Hall, London, 1965.

[5] J. Daemen, R. Govaerts and J. Vandewalle, Correlation matrices, in: *Fast Software Encryption 1994*, Lecture Notes in Comput. Sci. 1008, Springer, Berlin (1995), 275–285.

[6] J. Daemen and V. Rijmen, *The Design of Rijndael*, Springer, Berlin, 2002.

[7] G. R. Grimmett and D. R. Stirzaker, *Probability and Random Processes*, Oxford University Press, Oxford, 2001.

[8] X. Lai, J. L. Massey and S. Murphy, Markov ciphers and differential cryptanalysis, in: *Advances in Cryptology – EUROCRYPT 1991*, Lecture Notes in Comput. Sci. 547, Springer, Berlin (1991), 17–38.

[9] M. Matsui, Linear cryptanalysis for the DES cipher, in: *Advances in Cryptology – EUROCRYPT 1993*, Lecture Notes in Comput. Sci. 765, Springer, Berlin (1993), 386–397.

[10] S. Murphy, F. Piper, M. Walker and P. Wild, *Maximum likelihood estimation for block cipher keys*, Royal Holloway, University of London, technical report RHUL-MA-2006-3, 1994, www.ma.rhul.ac.uk/techreports.

[11] K. Nyberg, Linear approximation of block ciphers, in: *Advances in Cryptology – EUROCRYPT '94*, Lecture Notes in Comput. Sci. 950, Springer, Berlin (1995), 439–444.

[12] K. Nyberg, Correlation theorems in cryptanalysis, *Discrete Appl. Math.* **111** (2001), 177–188.

[13] S. D. Silvey, *Statistical Inference*, Chapman and Hall, London, 1975.

Author information

Sean Murphy, Information Security Group, Department of Mathematics,
Royal Holloway, University of London, Egham, Surrey TW20 0EX, UK.
E-mail: `s.murphy@rhul.ac.uk`