

# The Cocaine Auction Protocol: On The Power Of Anonymous Broadcast

Frank Stajano<sup>1 2</sup> and Ross Anderson<sup>1</sup>

<sup>1</sup> University of Cambridge Computer Laboratory,  
New Museums Site, Cambridge CB2 3QG, UK

<http://www.cl.cam.ac.uk/~fms27/>, <http://www.cl.cam.ac.uk/~rja14/>

<sup>2</sup> AT&T Laboratories Cambridge,  
24a Trumpington Street, Cambridge CB2 1QA, UK  
<http://www.uk.research.att.com/~fms/>

**Abstract.** Traditionally, cryptographic protocols are described as a sequence of steps, in each of which one principal sends a message to another. It is assumed that the fundamental communication primitive is necessarily one-to-one, so protocols addressing anonymity tend to resort to the composition of multiple elementary transmissions in order to frustrate traffic analysis.

This paper builds on a case study, of an anonymous auction between mistrustful principals with no trusted arbitrator, to introduce “anonymous broadcast” as a new protocol building block. This primitive is, in many interesting cases, a more accurate model of what actually happens during transmission. With certain restrictions it can give a particularly efficient implementation technique for many anonymity-related protocols.

## 1 Introduction

### 1.1 Why a cocaine auction?

Several extremely rich and ruthless men<sup>1</sup> are gathered around a table. An auction is about to be held in which one of them will offer his next shipment of cocaine to the highest bidder. The seller describes the merchandise and proposes a starting price. The others then bid increasing amounts until there are no bids for 30 consecutive seconds. At that point the seller declares the auction closed and arranges a secret appointment with the winner to deliver the goods.

Why are we interested in this scenario? One reason is that although electronic auctions are a very hot topic (witness the pioneering online auction house eBay [9], whose stock grew 1300% in only seven months after their IPO [19]), their privacy and trust implications have yet to be adequately discussed.

---

<sup>1</sup> To prevent grammatical rules from becoming sexist statements, one would normally split the roles of the principals between male and female personæ; but the drug dealers of our dramatisation look so much more plausible as ugly, cruel cigar-smoking men that it was felt more offensive than flattering to include any ladies. Of course some may now perversely call this an even more sexist statement...

In the eBay model, for example, the auction house's profit is a fixed percentage of the final sale, but at the same time bidders are asked to reveal to the auction house in confidence the maximum amount they are prepared to bid, so that the house can run the bidding process on their behalf without their having to be online for several days. This is putting the fox in charge of the hen house. The auction house could easily and undetectably exploit its knowledge of the bidders' limits to drive up the sale price—possibly introducing a fake bidder—in order to pocket the maximum commission. Users simply have to hope that the auction house will behave properly. EBay addresses some of the other trust concerns, such as whether users should trust other users: there is an interesting “peer review” system in which everyone gets a reliability rating from the principals with whom they interact. But while this mechanism may be valuable, it still cannot be used to justify the trustworthiness of the auction house itself.

We introduce the cocaine auction scenario as an exaggerated case that makes the trust issues unambiguous. We may assume that, in a game with such high stakes and shady players, nobody is going to trust anybody else any more than strictly necessary. We may also assume that the people who take part in the auction all know each other (otherwise one of them might be a police agent), but that no-one who places a bid wants to be identified to the other bidders or to the seller. Nobody except buyer and seller should know who won the auction; and even the seller should not be able to find out the identity of the highest bidder before committing to the sale. But none of the participants should have to trust any other: the protocol cannot rely on a judge or policeman to act as arbitrator and must instead be self-enforcing.

Data protection issues are a further reason for wanting an anonymous auction protocol. In the eBay model, each user has all her transactions logged, together with her personal data, and is at the mercy of abuses by the auction house and other participants, such as the resale of personal information to marketers, insurers or even thieves with a plausible tale (“May I please buy the list of all those who recently bought gold jewellery? I sell a really good polish.”). Serious political difficulties may arise from the US practice of storing and reselling such information as this violates the data protection laws of several other jurisdictions, including the European Union. If an auction site were to adopt protocols that prevented it from finding out the identity of bidders (or at least of unsuccessful bidders), then it would be in a much better position to claim that it had taken all reasonable steps to comply with data protection principles.

Finally, there was an amusing case of life imitating art when, the week before this paper was presented at the Information Hiding Workshop 1999, some Internet-savvy scoundrels *did actually offer a shipment of drugs on eBay*, with bids reaching 10 M\$ before the auction was noticed and shut down [18].

## 1.2 Anonymous broadcast

The second part of this paper examines the anonymity layer on which the auction protocol is built and proposes for it a provocative implementation technique

that does not use any cryptography. This novel approach offers substantial performance gains. Interestingly its security, while using mechanisms that were once considered dubious, turns out to be essentially equivalent to that of a cryptographically strong alternative, so long as we use realistic threat models.

Furthermore the anonymous broadcast primitive is also interesting from the protocol modelling point of view in that, for many cases, it gives a more faithful representation of what actually happens.

## 2 The cocaine auction protocol

### 2.1 Protocol

We shall now build our anonymous auction protocol assuming the availability of a mechanism for broadcasting messages to all the participants without revealing the identity of the sender (“anonymous broadcast”). The implementation of the anonymity layer will be discussed in section 3.

The other communication primitive we shall use is plain non-anonymous broadcast, where an identifiable principal delivers a message to all the others. Since in our scheme the identity of the seller is known to the buyers, the buyers’ messages are anonymous, but the seller’s are not.

The basic protocol is fairly simple and is organised as a succession of “rounds” of bidding. Round  $i$  starts with the seller announcing the bid price  $b_i$  for that round. Buyers have up to  $\Delta t$  seconds to make an offer (i.e. to say “yes”, meaning “I’m willing to buy at the current bid price  $b_i$ ”). As soon as one buyer anonymously says “yes”, he becomes the winner  $w_i$  of that round and a new round begins. If nobody says anything for  $\Delta t$  seconds, round  $i$  is concluded by timeout and the auction is won by the winner  $w_{i-1}$  of the previous round, if one exists. If the timeout occurs during round 0, this means that nobody made any offers at the initial price  $b_0$ , so there is no sale.

A couple of details need fixing before this protocol will work satisfactorily. Firstly, the seller needs a mechanism to identify the winner: if all the buyers ever say is just “yes”, then anybody can go to the seller, offer him the final sale price and obtain the goods in stead of the real winner—which is highly undesirable. This problem can be solved by ensuring that each “yes” message also contain a one-way function of a nonce: before completing the sale, the seller will then ask the winner to exhibit the original nonce, which nobody else could produce.

Secondly, once the auction is over, the seller might prefer to give a *secret* appointment to the winner (“see you on Tuesday at 06:30 in the car park of Heathrow terminal 2”) rather than exchanging suitcases of cocaine for cash under the noses of all the losing bidders. On the other hand, the identity of the winner should not be revealed to the seller until the latter commits to the sale. This is to protect the winner from the situation in which the seller says “So who won the auction? Oh, it was you? Well, anyone else would have been fine, but you’re from the wrong family, so I won’t sell to you after all, even if you’ve won”. To enable the seller to send an appointment to the winner only, but before knowing

the winner’s identity, we make the previously mentioned one-way function  $g^x \pmod n$ , where  $x$  is the nonce chosen by the bidder and  $g$  and  $n$  are public system-wide parameters. So each anonymous “yes” message will be of the form  $g^{x_i}$ , with  $x_i$  chosen arbitrarily by the winner  $w_i$  of round  $i$ . When the auction finishes, say at round  $f$ , with principal  $w_f$  committing to buy at price  $b_f$  with a “yes” message of  $g^{x_f}$ , the seller chooses a nonce  $y$  and performs a Diffie-Hellman key exchange [8] with the winner  $w_f$  (who is still anonymous) by broadcasting the appointment encrypted under the session key  $g^{x_f y}$ . The winner is the only buyer who can compute this key.

A number of minor variants to the protocol are possible. For example, before the first round the seller could specify, either algorithmically or by enumeration, the succession of bid prices  $\{b_i\}$ ; then he would no longer have to broadcast  $b_i$  at the beginning of each round, because each winner  $w_i$  would implicitly refer to the corresponding  $b_i$  in the “well known” succession.

A variant at the opposite end of the conciseness v. robustness trade-off [1] is to have the seller broadcast, at the beginning of round  $i$ , not only  $b_i$  but also the “yes” message ( $g^{x_{i-1}}$ ) of the winner of the previous round. This may help arbitrate races between bidders who both said “yes” to the same  $b_i$ . The bidders themselves could include in the “yes” message the value  $b_i$  to which they are responding.

We shall not, however, discuss this type of implementation detail any further; let us instead examine some of the ways in which the principals could misbehave, and whether the protocol is (or can be made) robust against them.

## 2.2 Attacks

There are limits to what can be achieved at the protocol level. It is always possible, for example, to subvert an auction when an appropriate combination of participants colludes against the others. For example, if all the bidders conspire against the seller, one of them can buy the shipment cheaply and divide it with the others later (a practice known as “ringing”), either in equal parts or perhaps even running a separate private auction and splitting the money that didn’t go to the seller. Similarly, if the seller plants an ally among the bidders, he can push the final selling price as high as the other bidders will bear (though at the risk of not being able to actually complete the sale). We do not believe that all such attacks can be detected, let alone stopped, by any particular auction protocol: the culprits can plausibly deny their involvement, since the “trace” of externally observable messages of an auction involving a collusion could always have been generated by a “honest” auction as well. Some specific conspiracies may be detectable by protocol level mechanisms, as we shall see; but we will not attempt to guard against the others.

**Seller not selling to highest bidder** Does the protocol force the seller to sell to the highest bidder? No, since the seller can always generate the session key starting with the  $g^x$  produced by whichever bidder he prefers, and nobody will

be able to tell that this happened just by looking at the commitment message. One might object that, so long as the price is strictly increasing from one bid to the next, the seller is guaranteed to lose money if he does not sell to the highest bidder—but the real world is more complicated than that, with all sorts of possible scams and double-crosses. We will differentiate two cases.

In the first, the seller sends an encrypted appointment message to the winning bidder but attempts some treachery in it. For example, he might have a sweetheart deal with one of the participants allowing him to match the best openly bid price, and might send the winner an appointment he has no intention to keep. In such a case, the winner's recourse is to blacken the seller's "good name" by complaining about the disappointment afterwards around the cocaine dealing community. We will consider this case no further.

In the second, the seller encrypts the appointment message using the  $g^x$  supplied by someone other than the winner. In this case, the cheated winner can broadcast an accusation and prove that the seller is dishonest, simply by publishing his last  $x$ . Anybody can then verify firstly that that  $x$  really corresponds to the  $g^x$  from the highest bidder and secondly that the message from the seller does not decrypt to anything meaningful using that  $x$ . At that point all the bidders take out their machine guns and the seller greatly regrets his dishonesty. Note that the cheated highest bidder manages to accuse the dishonest seller without exposing his identity, since he still sends out  $x$  anonymously. So this protocol ensures that most misbehaviour of this sort can be publicly exposed without loss of anonymity for the cheated accuser.

**Seller bidding at his own auction** The exception is of course where the seller bids at his own auction in order to raise the price, overshoots, and finds himself the winner. He is unlikely to accuse himself by broadcasting  $x$ ; but might he offer the shipment to anyone else?

Let us assume that he selects a  $g^x$  sent by one of the next-highest bidders, whom we will call Mr. N, and broadcasts an appointment message encrypted under  $g^{xy}$ , in the hope that Mr. N (whose identity is unknown to him) will keep quiet in exchange for a chance to buy the goods. When Mr. N sees that he can decrypt the secret appointment using his own  $g^x$ , he knows that the seller is cheating, since that message should have been sent to the highest bidder instead. So he can either expose the scam by exhibiting his own  $x$  as before and cause the barrels of all the Uzis in the room to converge on the seller; or he can accept the appointment and buy the shipment he would otherwise lose. (He might perhaps haggle over the price when he meets the seller there, but by then he will have lost his ability to cause immediate bodily harm to the seller from the comfort of anonymity).

So, when the seller tries to deal with someone other than the apparent winner, there seems always to be one principal who could expose him as a cheater, although in some sub-cases it is possible that the potential accuser might prefer to stay silent. A seller "with wife and kids", noticing that this situation carries a life risk greater than  $\varepsilon$ , might never attempt such a scam; but the more adven-

turous Scarface type on the fast track to the top might occasionally be prepared to run the risk.

Thus in practice the other principals still have no way of knowing for sure that an auction they lost was run “legitimately”. So they might want a way for all the participants to verify that the seller did encrypt the secret appointment to the highest bidder, although they should not be able to decrypt the message themselves. They might also want some reassurance that the appointment message decrypts to something meaningful. Both these assurances can be given by a cut-and-choose protocol. The seller broadcasts not one  $g^y$  but twenty (perhaps at the beginning of the auction) and once the auction concludes, he then offers a choice of twenty different encrypted appointment messages, such as “06:30 Tuesday in the car park of Heathrow terminal 2”, “23:20 Monday behind the George and Dragon”, ..., and will reveal up to nineteen of the  $y$  values in response to challenges. The same result might probably also be achieved using zero-knowledge proof techniques rather than cut and choose.

It may well be that our drug dealers do not care to pay for the extra complexity of such a protocol. If the main risk is felt to be the seller bidding at his own auction, then even if he overbids and is forced to sell to himself under penalty of being discovered, there is nothing to stop him from running a new auction some time later, pretending to have received a new shipment from his suppliers. (This applies to commodities like cocaine or memory chips, but not to unique and recognisable items such as stolen Rembrandts, where a cut-and-choose protocol might be preferred.)

**Deadbeat bidders** A general consequence of anonymity is that it is hard to hold anonymous principals responsible for anything. In particular, it is hard to guard against “deadbeat bidders”, i.e. participants who win an auction and then do not turn up with the cash to buy the goods. With the protocol described so far, they would get the encrypted appointment but not show up, and the seller would not know who to blame. While deadbeat bidders would not gain much (and would certainly not get any free cocaine), their behaviour will certainly annoy the other participants. If repeated undetectably over a series of auctions, deadbeat bidding could amount to a denial of service. One might argue that nonpayment lies outside the scope of the auction protocol, which should only designate a winner and a sale price; but it is still reasonable to seek ways in we might at least identify persistent deadbeats.

The approach used by “respectable” online auctioneers such as eBay is to enable clients to build up a reputation for honest dealing. One might try to transplant this to the cocaine auction by giving each principal a pseudonym; but as soon as the winner turns up to collect his cocaine, the link between his name and pseudonym becomes obvious, unless the actual delivery of goods is also conducted anonymously. In fact, even in the “respectable” case, service denial attacks can be mounted by any principals who can repeatedly acquire new identities.

This problem raises complex issues related to identity certification, which in itself might be a concept that our mistrustful drug dealers unconditionally reject *a priori*. Here we will merely point out a major practical pitfall. Suppose that some acceptable technical mechanism has been devised to create a certification authority. For example, one might set up a  $k$ -out-of- $n$  identity escrow scheme with a very high  $k$ , say  $3n/4$ : after a certain level of missed appointments, or deadbeat bids, were detected by the audience, everybody could cooperate to unmask the disruptor, in a way reminiscent of Blaze’s “angry mob cryptanalysis” [4]. Mechanisms for setting up the underlying threshold signature schemes without a trusted party are known [7]. But the real problems are likely to come from the application detail, such as the plausibility of the excuse that the subject might put forward to justify his absence (was he arrested “honestly”, or did he pay the police to arrest him?).

**Do auction houses have a future?** One of the questions asked by many businesses is whether the convergence of computers and communications could destroy their niche. Even banks worry about “disintermediation” as their corporate customers raise loan capital on the markets directly. What does the future hold for auctioneers?

A traditional auction house adds value in a number of ways. Some of these, such as marketing, may be easier to do on the net; others, such as providing assurance that the goods on offer are as described and are the lawful property of the seller, may be harder. But comparing the cocaine auction with existing ones does tell us something about transaction costs. In conventional auctions, bidders must identify themselves to the auctioneer, who can exclude any known deadbeats; and although there is no formal mechanism to detect when a friend of the seller is bidding secretly for him, there is a deterrent in that a seller who buys his own merchandise ends up out of pocket by the auctioneer’s commission.

In many sectors of the economy, from securities trading to the provision of airport taxi services, it has been found that regulated markets build confidence and attract custom. The question for regulators is precisely how much to regulate. We hope that comparing conventional auctioneers (and the new online firms such as eBay) with the fully disintermediated environment of the cocaine auction protocol may provide some useful insight.

### 3 The anonymity layer

To make the cocaine protocol usable, we must also supply a mechanism that allows the bidders to anonymously broadcast their “yes” message.

#### 3.1 The dining cryptographers

The “dining cryptographers” construction introduced by Chaum [6] addresses precisely this problem. In his now classic story, several cryptographers are gathered around a table for dinner, and the waiter informs them that the meal has

already been paid for by an anonymous benefactor, who could be one of the participants or the NSA. The cryptographers would like to know whether they are indebted to one of their own number or to the agency. So Chaum incrementally constructs a protocol through which, after the sharing of what are effectively one time pads between selected pairs of principals, each principal outputs a function of her “I paid/I didn’t pay” bit and everyone can later work out the total parity of all such bits. As long as not more than one of the cryptographers says “I paid”, even parity means that the NSA paid, while odd parity means that one of the diners paid, even if nobody can figure out who.

Various extensions are then proposed, including one in which the principals are arranged in a token ring and transmit integrity-checked blocks rather than single bits, so that collisions<sup>2</sup> can be detected.

For the system to work, collisions must be controlled by ensuring that the round trip period of the token is much smaller than the typical time between the seller announcing a new price and a bidder responding “I’ll buy”. Furthermore, the procedure for dealing with collisions must take into account the nature and aims of the principals: it might be inappropriate for us to simply invite colliders to retransmit after random intervals, as Chaum does, since two principals might both retransmit as soon and as often as possible in an attempt to secure the bid, thereby causing collisions *ad infinitum*.

Pfitzmann proposes an ingenious optimisation [16] that guarantees to resolve a collision between  $n$  participants in at most  $n$  rounds—while the probabilistic approach suggested by Chaum cannot guarantee an upper bound.

Despite this, the cocaine auction protocol implemented using “token ring dining cryptographers” is fairly expensive in terms of communications. Even ignoring the initial cost of setting up the pairwise one time pads and of any retransmissions caused by collisions, each participant must send at least one message to his neighbour for each round in which one untraceable bid may be sent by one of the parties, plus another one for the second trip round the loop which is needed to communicate the result to all participants. Calling  $n$  the number of participants,  $r$  the number of rounds of bidding needed before the auction results in a sale and  $K$  the “dilution factor” introduced to spread out transmissions over several rounds so as to minimise collisions, the protocol so far described requires  $2 \cdot n \cdot r \cdot K$  such messages to be sent. (Pfitzmann’s cited construction allows  $K$  to be much smaller than in Chaum’s case by lowering the retransmission costs on collisions.)

We shall now show how a simple assumption about the nature and properties of the physical transport layer, directly inspired by a specific implementation technology, dramatically reduces these costly transmission requirements.

---

<sup>2</sup> Only one cryptographer at a time can say “1” and be heard. Since the resulting anonymous bit is the XOR of those transmitted by the individual principals, if at most one principal transmits a “1” and all others transmit a “0”, we can infer as appropriate that nobody said “1” or that one anonymous principal said “1”. But if more than one principal sends a “1”, the protocol no longer works—it can only tell whether an even or an odd number of principals transmitted a “1”.



### 3.2 Anonymous broadcast based on physics

The original idea for the cocaine auction protocol arose in the context of the discussion about possible future applications for a short-range radio networking facility such as that provided by Piconet [3]. We envisage that the drug dealers of our story might hold in their pockets little radio transmitters similar in shape and size to car key fobs, and that unobtrusively pressing the button on the device would cause the transmission of the relevant “yes” message (a transmitter with a slow processor, unable to do modular arithmetic quickly, might rely on a precomputed list of  $g^x$  for various  $x$ ).

By using radio, each message sent is automatically broadcast to all principals; and it can be anonymous, as long as we simply omit to mention the sender in the link-layer header. Only one such message is needed per auction round, so in terms of transmissions the entire auction only costs  $r$  messages, as opposed to  $2 \cdot n \cdot r \cdot K$  (plus the extras we hinted at) for the dining cryptographers implementation.

The savings are dramatic and worth investigating in greater detail. As it turns out, they come from having questioned a basic assumption in protocol modelling, namely that communication is point-to-point. This trick can be exploited in a variety of interesting cases that have nothing to do with auctions.

### 3.3 A fundamental protocol building block

Traditionally, cryptographic protocols are described as a sequence of steps of the form

$$A \rightarrow B : M$$

indicating that principal  $A$  sends message  $M$  to principal  $B$ . In the general case it is proper to assume a primitive that sends a message from a specific sender to a specific recipient; indeed in most cases this is what the communications API offers. Other constructions are typically derived from this one: for example, broadcasting the same message  $M$  to all the members of a domain  $\mathcal{D}$  can be represented by a trivial, if inefficient, iteration. Anonymous sending can be achieved by more elaborate constructions which rely on building an untraceable tortuous path across a multiply connected network, in which intermediate relay nodes hide the route that the message is going to take next [5,10], or on diluting the payload in a cloud of messages sent by a community of principals, each of which might have been the actual sender [6], or on broadcasting a message which only some subset can decipher [4].

Let us now reverse the perspective and take “anonymous broadcast” as the fundamental primitive. The following notation

$$A \text{ ? } \longleftrightarrow \mathcal{D} : M$$

shall signify that principal  $A$  broadcasts message  $M$  anonymously into the domain  $\mathcal{D}$ . This means that all principals in  $\mathcal{D}$  receive  $M$ , and given any two messages  $M_1$  and  $M_2$  that have been anonymously broadcast at different times into  $\mathcal{D}$ , no principal in  $\mathcal{D}$  (except the originators) is able to tell whether they came from the same principal or not.

From here we can derive non-anonymous broadcast by naming the sender in the message, as in

$$A \xleftrightarrow{\mathcal{D}} M \equiv A ? \xleftrightarrow{\mathcal{D}} (A, M)$$

and point-to-point send by naming both the sender and the recipient, as in

$$A \rightarrow B : M \equiv A ? \xleftrightarrow{\mathcal{D}} (A, B, M)$$

Many security papers are concerned with the strength of the mechanisms used to bind names such as  $A$  and  $B$  to  $M$ : but we are not concerned with non-repudiation, only with its dual, namely plausible deniability. We simply note that a basically anonymous broadcast primitive, coupled with weak mechanisms to name the claimed sender and the supposedly intended receiver, is what *really* happens in practice in many common cases, including radio and ethernet. (Human speech is also a form of local broadcast, anonymous only in crowds, and yet there are many places—not just Japanese paper houses—where it’s conventional to ignore messages addressed to others.)

At the physical level, whether in the one-bit bus of ethernet or in the “ether” of radio waves, it is actually the point-to-point messaging facility that is obtained as a composite construction built on top of anonymous broadcast. The practice of prefixing each outgoing message with the globally unique ID that ethernet adapters receive during manufacture is a convention which a malicious node can easily ignore.

Physical anonymous broadcast as a genuine networking primitive requires a shared communication medium and thus is only practical in local networks. In the wide area, its usefulness is limited by issues of transmission power, propagation delay and bandwidth: if only one principal can transmit at a time without collision, larger domains mean that more principals are forced to stay silent while one of them is transmitting. Above a certain domain size, routing becomes the preferred option in order to reduce the transmission power and increase the aggregate bandwidth. So basing a cryptographic protocol on the anonymity properties of the physical broadcast primitive limits it to the local area. The drug barons can hold the efficient version of the auction around a table, with key fobs in their pockets or with an ethernet connecting their laptops, but to hold an auction over the Internet from their swimming pools it appears that they would have to implement the anonymity layer using the more tedious classical techniques. This problem is also faced by Jackson [11], who finds he cannot use the otherwise desirable physical broadcast technique in an active badge application because of power and scalability limitations.

In summary, the anonymous broadcast primitive has two main advantages. The first is efficiency: even before taking anonymity into consideration, in a shared transmission medium sending one message from  $A$  to  $B$  or sending one message from  $A$  to anyone else in the domain has exactly the same cost. It is foolish to use a formalism that hides this, and end up being forced to send  $\#\mathcal{D} - 1$  messages rather than one when broadcast is really intended. Moreover, under some assumptions that we shall examine in greater detail next, a shared medium can give anonymity practically for free, thus saving the many redundant

messages otherwise needed to hide the real traffic. It is thus convenient to be able to leverage off these valuable properties when designing a higher-level protocol.

The second reason is clarity: by more closely modelling what goes on during the transmission of a message, we are in a better position to evaluate the actual security properties of our system.

### 3.4 The strength (or weakness) of broadcast anonymity

Although the network addresses sent out by most nodes in conventional computer and communications systems can be forged easily, it would be naïve to conclude that every transmission on a shared physical medium provides strong anonymity. Pfizmann, who extensively analysed techniques to reduce user observability in communication networks [13,14,15,16,17], mentioned broadcast over a shared medium as a possible low-level anonymity layer, but then dismissed it in favour of more elaborate solutions based on cryptography as it could not provide unconditional security. This is a valid objection: an opponent who can observe the transmissions at a low enough level will generally be able to distinguish the participants.

For example, in cable and radio transmissions, at the physical level the distance between transmitter and receiver affects both the received power and the transmission delay. Several conspiring nodes who could observe these parameters might well be able to identify the senders of most of the messages. We are not interested in the trivial case in which all nodes but one conspire against the last, but in more serious cases where three or four nodes can, between them, tell whether the last message came from here or from there.

A transmitter might randomly vary its power output to prevent direction finding using signal strength measurements; but the relative amplitude will still be visible, and precise timing information will anyway be what tells most to a well equipped attacker. The game could include directional receiving and transmitting antennas; a “defensive” environment such as a naked Faraday cage that maximises reflections and echoes; or a “hostile” environment with unseen receivers in every ceiling tile and surveillance cameras in every pot plant. Bidders might even use extreme electronic warfare techniques to try to frame each other by causing certain messages to be received by only some subset of participants. Radio-frequency ID systems can identify individual transmitters by their analogue characteristics; to block this, the principals might have several buttons in their pockets which they use once each as they wander around.

However, many of the attacks open to a resourceful opponent are independent of the strength of cryptography in use. Equipment can be abused; keyboard sniffers can be installed in laptops and signal using soft tempest transmitters [12]. If this is impossible, cameras and microphones can be used to observe user input. There may be infrared lasers to measure bidders’ blood oxygen levels and special toilets to collect biological samples to measure stress. One should bear all this in mind when dismissing anonymous broadcast as insecure compared with the “mathematically provable” untraceability of the dining drug dealers’ scheme. An opponent who forces us to think about unconditionally secure cryptography may

also fly below the mathematics and steal the plaintext at the level of electronics, physics or biology. In such cases, the strength of cryptographic mechanisms is only one factor of many. It is probably fair to say that the attacker equipped to break the anonymity of our physical broadcast will have little trouble in bypassing strong cryptography (wherever it is employed) by simply observing the keys or plaintext directly.

At less exalted levels, we have to consider the economics of the attacker and the defender. It is highly significant that defence (enabling anonymous broadcast) may require only a small change in the firmware of a standard node, and might even be performed in software depending on how “shallow” the network adapter is; but attack (sensing the signal strength, timing, phase, polarisation and/or analogue transmitter characteristics of received messages) is a lower level and thus more expensive operation.

## 4 Conclusions

We have presented a protocol for anonymous auctions among mutually mistrustful participants and examined its vulnerability to various attacks. Drug dealers made convincing actors for our dramatisation, but our analysis is more general. Online auction houses, commercial companies seeking tenders and even governments calling for bids might all benefit from a scheme in which bidders do not have to reveal their strategies to each other, and in which unsuccessful bidders remain anonymous. As online auctions gain importance and even governments try to move their business online, so will the issues of anonymity, privacy and trust become more critical.

Our analysis showed that introducing a realistic non-cryptographic assumption about the physical nature of the communications medium gave us a dramatic efficiency gain. This “anonymous broadcast” assumption maps directly to the implementation of local area networking environments based on a shared transmission medium such as ethernet or radio. While the strength of its anonymity depends on the threat model, the kind of attacks that might defeat it are uneconomic in many applications and are also likely to defeat “provably secure” cryptography. So, where appropriate, anonymous broadcast may offer significantly better performance, as well as a more accurate model of what is actually going on.

The anonymous broadcast technique is not applicable to the wide area case, so the more efficient version of our auction protocol can only be run locally. But the issues we raised about attacks and deceptions are orthogonal to how the anonymity is implemented: our discussion may provide a useful yardstick even for online auction implementations that use a different anonymity primitive or a different auction protocol.

We conclude by observing that new ideas come from challenging fossilised axioms: it is needlessly limiting to design all cryptographic protocols under the unexamined assumption that communications must perforce be point-to-point.

## 5 Acknowledgements

We thank the audiences that gave us their valuable feedback during presentations of this research: in particular we appreciated the comments of Bruce Christianson, Markus Kuhn, Stewart Lee, Roger Needham, Larry Paulson and Mike Roe from the security group at the University of Cambridge; of John McHugh and Nicko van Someren at the Information Hiding Workshop in Dresden; and of Raymond Liao at Columbia University. Further useful remarks were offered offline by Andreas Pfitzmann and George Danezis.

## References

1. Martín Abadi and Roger Needham. Prudent engineering practice for cryptographic protocols. Technical Report 125, Digital Equipment Corporation Systems Research Center, June 1994. <ftp://ftp.digital.com/pub/DEC/SRC/research-reports/SRC-125.pdf>.
2. Ross Anderson, editor. *Information Hiding: First International Workshop proceedings*, volume 1174 of *Lecture Notes in Computer Science*. Springer-Verlag, 1996.
3. Frazer Bennett, David Clarke, Joseph B. Evans, Andy Hopper, Alan Jones, and David Leask. Piconet: Embedded mobile networking. *IEEE Personal Communications*, 4(5):8–15, October 1997. <ftp://ftp.uk.research.att.com/pub/docs/att/tr.97.9.pdf>.
4. Matt Blaze. Oblivious key escrow. In Anderson [2], pages 335–343. <http://www.crypto.com/papers/netescrow.ps>.
5. David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, February 1981. Unofficial copy at <http://www.wiwi.uni-frankfurt.de/~kcotoaga/offline/chaum-acm-1981.html>.
6. David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1:65–75, 1988. Unofficial copy at <http://www.scu.edu/SCU/Programs/HighTechLaw/courses/ccp/diningcr.html>.
7. Clifford Cocks. Split knowledge generation of RSA parameters. In Mike Darnell, editor, *Cryptography and coding: 6th IMA conference, Cirencester, UK, December 17–19, 1997: proceedings*, volume 1355 of *Lecture Notes in Computer Science*, pages 89–95. Springer-Verlag, 1997. <http://www.cesg.gov.uk/downloads/math/rsa.pdf>.
8. Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Trans. Inform. Theory*, IT-22(6):644–654, November 1976.
9. eBay. <http://www.ebay.com/>.
10. David M. Goldschlag, Michael G. Reed, and Paul F. Syverson. Hiding routing information. In Anderson [2], pages 137–150. <http://www.onion-router.net/Publications/IH-1996.ps>.
11. Ian W. Jackson. *Who goes here? Confidentiality of location through anonymity*. PhD thesis, University of Cambridge, February 1998. <http://www.chiark.greenend.org.uk/~ijackson/thesis/>.
12. Markus G. Kuhn and Ross J. Anderson. Soft tempest: Hidden data transmission using electromagnetic emanations. In David Aucsmith, editor, *Information Hiding: Second International Workshop*, volume 1525 of *Lecture Notes in Computer Science*, pages 124–142. Springer-Verlag, 1998. <http://www.cl.cam.ac.uk/~mgk25/ih98-tempest.pdf>.

13. Andreas Pfitzmann. Ein dienstintegriertes digitales Vermittlungs-/Verteilnetz zur Erhöhung des Datenschutzes (*An Integrated Digital Services Switching/Distribution Network for Increased Privacy*). Technical Report 18/83, Institut für Informatik IV, University of Karlsruhe, 1983.
14. Andreas Pfitzmann. A switched/broadcast ISDN to decrease user observability. 1984 International Zurich Seminar on Digital Communications, Applications of Source Coding, Channel Coding and Secrecy Coding, March 6-8, 1984, Zurich, Switzerland, Swiss Federal Institute of Technology, Proceedings IEEE Catalog no. 84CH1998-4, 6-8 March 1984.
15. Andreas Pfitzmann. How to implement ISDNs without user observability—some remarks. Technical report, Institut für Informatik, University of Karlsruhe, 1985.
16. Andreas Pfitzmann. *Diensteintegrierende Kommunikationsnetze mit teilnehmerüberprüfbarem Datenschutz* (Integrated services communication networks with end-user verifiable privacy). Number 234 in Informatik-Fachberichte. Springer-Verlag, Heidelberg, 1990.
17. Andreas Pfitzmann and Michael Waidner. Networks without user observability. *Computers and Security*, 6(2):158–166, April 1987. [http://www.semper.org/sirene/publ/PfWa\\_86anonyNetze.html](http://www.semper.org/sirene/publ/PfWa_86anonyNetze.html).
18. Greg Sandoval. eBay auction goes up in smoke. *CNET*, September 1999. <http://news.cnet.com/news/0-1007-202-123002.html>.
19. StockMaster. <http://www.stockmaster.com/exe/sm/chart?Symbol=EBAY>.