Megamos Crypto, Responsible Disclosure, and the Chilling Effect of *Volkswagen Aktiengesellschaft vs Garcia, et al*

Robert Carolina and Kenneth G. Paterson

August 28th 2013

1. Introduction

The recent decision of the English High Court to censor the publication of an academic paper describing weaknesses in the Megamos Crypto automobile immobiliser system raises a number of concerns for members of the cryptographic academic community, legal practitioners, and commercial users of cryptographic products.

In this paper we will provide a brief description of the technology at the heart of the dispute, the crypto research project, the court's decision, and then provide a critique of the decision and make observations about its potential impact.

Our description and our observations are based on evidence as it was disclosed in the published decision of the court. *Volkswagen Aktiengesellschaft vs Garcia, et al*, [2013] EWHC 1832 (Ch) (25 June 2013). This decision addressed a request for preliminary injunction pending a full trial on the merits. It remains possible that additional evidence introduced later, or existing evidence that has not been disclosed in the decision, could have a significant impact upon the observations and opinions presented here. We do not take any position, nor do we make any prediction, about the ultimate outcome of this case.

2. The Megamos Crypto immobiliser system

The case arises from the desire of three academic researchers who wish to discuss potential weaknesses they have discovered in a cryptography-based automobile immobiliser system known as Megamos Crypto. This system (in common with many others like it) is designed to act as a deterrent to automobile theft by placing the automobile in an immobilised state until a properly coded hand-held device is brought within close proximity of the vehicle.

From the description provided by the court, Megamos Crypto is a symmetric crypto system. The immobiliser device embedded in the automobile is paired with a hand-held device in the driver's possession. These paired devices hold a shared secret key that is supposed to be unique for each protected

automobile (or possibly unique to each hand-held device). Both devices include a specialist microprocessor chip incorporating a cryptographic algorithm that is common to all of the Megamos devices.

The immobiliser system operates on a challenge-and-response basis. The embedded device generates and transmits a random number to the hand-held device. The random number and the shared secret key number constitute two inputs to the algorithm that calculates a resulting third number. (Para.5.)

The calculated number is then split into two halves. The device in the car transmits one half of the calculated number to the hand-held device. The hand-held device compares this with its own calculation. This is how the device in the car proves (authenticates) its identity to the hand-held device. Having been persuaded that it is talking to the correct car (and not a malicious third party), the hand-held device transmits the second half of the calculated number. This is used to authenticate the identity of the hand-held device to the embedded immobiliser. Once the embedded device is convinced that a paired hand-held device is present, the immobiliser de-activates and it becomes possible to operate the car. (Para.5.)

The supply chain for the crypto immobiliser product, in common with many aspects of automobile manufacture, is long. Thales created the crypto algorithm. (Thales will be joined into the lawsuit as an interested party.) Thales authorised a company called EM to manufacture microprocessor chips that incorporate the algorithm. Each chip incorporates the algorithm. EM sells these chips to Delphi. Delphi installs the chips into immobiliser hardware. Delphi then sells the complete immobiliser package to Volkswagen, and other automobile manufacturers, who install these in automobiles during the assembly process. (Para.7.) Volkswagen is said to have installed this immobiliser in "millions" of automobiles. (Paras. 2, 18, 28, and 44.)

3. The research project, the private disclosure, and the lawsuit

The three researchers decided to test the strength of the Megamos Crypto system. This type of activity – an unsolicited effort to identify weaknesses in commercial crypto devices – is common in the field of crypto research. This would not be the first paper published by academics highlighting weaknesses in RF-based automobile security devices. See, for example, Indesteege, Keller, Dunkelman, Biham, and Preneel, "How To Steal Cars – A Practical Attack on KeeLoq", EUROCRYPT 2008, LNCS 4965, pp. 1–18, 2008. (Cooperative effort by academic researchers resident in Belgium and Israel, supported by both public and private research grants, revealing deficiencies in KeeLoq – a widely installed remote key entry system. We note without further comment that it is common practice for academics in this field to give such papers rather provocative titles.)

To conduct their analysis of Megamos Crypto the researchers needed to obtain details of the crypto algorithm used. The manufacturers of the immobiliser do not publish the algorithm. The algorithm is claimed as a trade secret. It is not clear from the decision whether the researchers considered

paying a laboratory to reverse engineer the crypto chip itself. The court was advised that reverse engineering the chip would cost in the region of €50,000 – an outlay that might have seemed expensive in the context of an academic grant proposal. Instead, the researchers identified a third party hardware and software product called Tango Programmer. This product (sold for an initial payment of €1,000 per unit) can be used, among other things, to create keys for automobile immobilisers using Megamos Crypto and other immobiliser systems. The algorithm is incorporated within the Tango Programmer software, but not directly disclosed by the manufacturer.

The researchers conducted a careful study of the Tango Programmer software. From this they were able to reverse engineer the functionality of the system and discover the details of the cryptographic algorithm. Having obtained the algorithm, they set about to study the system.

The researchers eventually identified three weaknesses in the system. (Para.11.) Two of these (use of weak secret keys and poor key updating practices) were not an issue in the case. The court did not comment on this, but we note that weaknesses of this type recur with sad frequency in the operation of secure systems.

The other weakness identified is much more serious. This is alleged to be a weakness in the design of the cryptographic algorithm itself. To explain the flaw that they had uncovered, the researchers planned to include a description of the algorithm in their published paper. It was this desire to publish the (allegedly secret) algorithm that created the dispute.

In November 2012 the three authors approached EM (the crypto chip manufacturer) to explain the weaknesses they had uncovered. (Para.15.) It is not clear from the decision whether the researchers understood that EM was using the algorithm under license from Thales. The court's published decision, unsurprisingly, does not provide details of the exact nature of the weakness in the Megamos Crypto algorithm.

The researchers planned to publish their paper in August 2013 in the proceedings of the annual USENIX Security conference. Volkswagen learned of the soon-to-be published academic paper on 23 May 2013, and brought a lawsuit in the High Court of England to prohibit disclosure of the algorithm. (Para.16.) The lawsuit names the three academic authors (two resident in the Netherlands and one resident in England), and the two universities that employ them (in England and the Netherlands).

4. The relevant law

The law of trade secrets and confidential information provides relief only to people whose confidential information is improperly taken. In the context of industrial trade secrets it is generally accepted in the courts of England, Commonwealth countries, and the United States, that there is a critical distinction between obtaining secrets by improper means (which includes any number of unsavoury acts such as bribing a member of staff to disclose the

secret) and reverse engineering (discovering a secret through the careful study of a lawfully obtained product). Once a secret has been lawfully discovered and made generally known it ceases to be confidential.

The researchers discovered details of the algorithm through reverse engineering, but not by reverse engineering the Megamos Crypto product itself. They instead reverse engineered third party software (Tango Programmer) that makes use of the crypto algorithm.

But how did the allegedly secret algorithm find its way into the Tango Programmer product? If Tango Programmer had been created using a properly reverse-engineered trade secret, then Tango Programmer would not infringe rights under the law of confidential information. If, on the other hand, Tango Programmer was created using a misappropriated trade secret this makes Tango Programmer an infringing product and exposes those who manufacture sell or use the product to potential liability. For the researchers, a critical question is then whether they actually knew – or should be deemed responsible to have inferred – the presence of a misappropriated trade secret. Using the rather fanciful language of English case law, to remain free from liability the person receiving such trade secret information cannot escape liability if they have had their conscience affected.

Because the requested court order would serve to restrain publication, the court was required to analyse the request under the heightened scrutiny of freedom of speech law. The standard used by English courts to decide such prior restraint cases comes from the Human Rights Act. This law says that a preliminary injunction to restrain publication can be granted only if the court believes that a full trial is "likely to establish that publication should not be allowed". (Para.24, quoting the Human Rights Act s.12(3).)

The House of Lords (at that time the supreme judicial authority of England) provided guidance to the interpretation of this standard in *Cream Holdings v. Banerjee*, [2004] UKHL 44. The judgment in that case delivered by Lord Nicholls explained that a party seeking a preliminary injunction to restrain publication must demonstrate a case "sufficiently favourable to justify such an order being made *in the particular circumstances of the case*" (emphasis added). Although this usually means that a court can only restrain publication if it believes that it is more likely than not that the complaining party will win, this is not an absolute rule. A lower probability of success is acceptable "where the potential adverse consequences of disclosure are particularly grave, or where a short-lived injunction is needed to enable the court to hear and give proper consideration to an application for interim relief pending the trial or any relevant appeal". (*Cream Holdings v Banerjee*, op. cit., para.22.)

5. The court's decision

The court concluded without much discussion that a full trial would probably establish that the algorithm remains a trade secret belonging to Thales. (Para.27.) Although the court seems highly confident in this conclusion, the decision sheds relatively little light on why. The court seems to accept that

some third party somewhere has successfully reverse engineered the algorithm, although apparently without publishing the full result. The decision also makes clear that the algorithm has been incorporated into a version of Tango Programmer that has been openly available for sale since 2009. Notwithstanding the rather broad group of people who now hold details of this algorithm, the court concluded that it continues to have "the necessary quality of confidence" required for legal protection.

A significant amount of the court's decision is taken up with the question of how the algorithm came to appear within Tango Programmer, and what the researchers should have known about that. The decision reveals surprisingly little evidence on whether the algorithm was obtained by the manufacturer of Tango Programmer through reverse engineering a secret (which does not violate rights in confidential information), or misappropriating a secret (which does).

The court accepted that it is indeed possible to reverse engineer the algorithm from the microchip embedded in the product. The method discussed by the court is known as chip slicing: studying the microchip one slice at a time using an electron microscope. The court was presented with evidence that this could be done by a suitably equipped laboratory at a cost of €50,000. In fact, the court was presented with and accepted evidence that other people have already chip sliced Megamos Crypto. The decision mentions "Silicon Zoo" by name, but notes that there was no evidence presented that anyone who has conducted this analysis has published the algorithm as a result. (Para.36)

The court makes clear that it had serious suspicions about the manufacturer of Tango Programmer. "I think it is obvious that Tango Programmer does not derive from a legitimate source in the automotive industry." (Para.36.)

When parties to a case concerning alleged violation of rights in confidential information (like the researchers) do not themselves stand accused of breaching their own direct obligation of confidentiality, but instead are accused of misusing a secret received from someone who has breached such an obligation, it is important to establish what the recipient knew and when they knew it. Having reached the preliminary view that Tango Programmer was created using a misappropriated trade secret, the court turned its attention to what the academic researchers knew about this state of affairs when they obtained and then reverse engineered Tango Programmer.

The court was not restrained it is criticism. "… Tango Programmer has a clearly murky origin, and that is obvious to the [academics]. Despite this being obvious to them, the [academics] have made no effort to make enquiries about the legitimacy of Tango Programmer. That is not to impose a (sic) unreasonable burden on them. They are the ones who obtained the information from Tango Programmer. [One author who gave a witness statement to the court] simply asserts that it must have been chip sliced. I do not accept that. In my judgment, the [academics] have taken a reckless attitude to the probity of the source of the information they wish to publish." (Para.38.)

Applying the rule from *Cream Holdings*, the court decided that Volkswagen's case (on the limited evidence available) was "much more than … merely arguable … [but] not overwhelming". While hardly an endorsement of the chances of success at trial, the judge concluded there is sufficient evidence to justify at least a temporary injunction until a trial can take place. (Para.43) The court went on to stress that this (for the moment temporary) infringement of free speech and academic freedom had been balanced against "the security of millions of Volkswagen cars". (Para.44). This suggests that the court also relied on the second limb of the rule from *Cream Holdings.* The court rather clearly believed that the potentially adverse consequences of publishing the algorithm would be "particularly grave".

The court did not forbid publication of the entire academic paper. It ordered that certain parts of the paper should be withheld from publication until a trial can take place. The court's published decision does not (of course) disclose any technical detail about the content to be withheld.

6. Observations and Criticism

There is much that we find troubling about the decision itself. In this section we will discuss: (1) the inherent problem of attributing value to a "secret" crypto algorithm; (2) the lack of analysis concerning the quantum of risk attending publication of this paper; (3) the court's willingness to infer that the algorithm was discovered through misappropriation, and the burden of proof on this question; (4) the state of knowledge required to demonstrate liability against a third party accused of misusing a secret that was misappropriated by another person; (5) the apparent confusion surrounding the meaning and purpose of "responsible disclosure" when used as a term of art in security research; and (6) the apparently long delay of the complaining parties in enforcing their rights in the alleged trade secret.

6.1 The value of a "secret" crypto algorithm

The court makes a surprising statement early in the decision: "For the process to be secure, both pieces of information need to remain secret - the key and the algorithm." (Para.5) To a cryptographer, this claim is puzzling. It is a well-known and widely accepted maxim in the field of cryptographic system design that such systems should remain secure even when the crypto algorithm falls into the hands of a malicious third party. The strength of cryptographic systems instead depends on the idea that it is infeasible (but not impossible) to use technological means to decrypt or forge messages or to discover a secret key, *assuming that the algorithm is freely available for study*. This is why it is not technically possible to conduct a study of cryptographic strength without access to the algorithm, and why the researchers wanted to procure these details.

Although cryptographic systems are in general supposed to remain secure when the algorithm is known, some security product manufacturers attempt to maintain the secrecy of their algorithm. While manufacturers are not

prohibited from attempting to keep an algorithm secret, if the publication of that algorithm is completely devastating to the strength of the system a cryptographer can only conclude that the algorithm itself is seriously flawed. In other words, publication of an algorithm only weakens a system to the extent that the algorithm itself fails in its core purpose.

Although the complaining parties are working hard to support the proposition that the Megamos Crypto algorithm remains a secret, like most crypto algorithms incorporated into a massively distributed consumer product it is – at best – a secret that is poorly kept and unlikely to remain secret indefinitely. In this case, the court essentially acknowledged that the algorithm – alleged to be critical to the security of millions of automobiles – could be (and in fact has been) legally discovered, and thus could be legally disclosed at any time, by anyone who has already discovered it or anyone else who is willing to pay €50,000 to replicate the reverse engineering work.

In our opinion, the lack of understanding of the distinction in maintaining the secrecy of the algorithm (common to all devices) and secrecy of the keys (unique to each automobile) demonstrated in the decision substantially undermines the remaining analysis and thereby the credibility of the decision as a whole.

6.2 Quantifying the risk of disclosure

As a prior restraint of speech case, the court made clear that this decision rested in part on its conclusion that publication would put the security of millions of cars in jeopardy. (Para.28) In doing so, the court makes a qualitative observation about the impact of publication (i.e., it will decrease security). But it does not present a quantitative statement about the impact of publication (i.e., to what degree it will decrease security). Rather than grappling with this risk analysis, the court was prepared to conclude this analysis with the simple observation that somebody somewhere will eventually exploit this weakness. Even if one accepts this, it does not necessarily lead to a conclusion that the consequences of publishing the secret algorithm would be "particularly grave".

There is at least one counter-argument to the court's conclusion that the consequences of publishing the algorithm would be particularly grave that was not addressed in the decision and helps to shed some light on the role of a quantitative risk analysis. If access to the algorithm alone changed the economics of theft from unreasonably expensive to trivially cheap, and it is possible to reverse engineer the algorithm for the sum of €50,000, then there can be little doubt that some enterprising criminal gang (or gangs) would have done so already. The input cost (which looks excessive to a university academic) could be recovered with the theft of even one high-priced automobile. If a criminal gang has not done this already, it suggests that the algorithm alone does not have such an enormous impact on the economics of criminal risk-taking and therefore the risk of publishing the algorithm is not that high. If a criminal gang has in fact already done this, then the secret is already

in the hands of the criminal underworld and – once again – the risk to society of further disclosure is not that significant.

This issue of quantifying risk leads to another possibility that the court seems to have rejected – the risk inherent in delaying publication. The court seems to proceed on the (hopeful) assumption that these researchers are the first people to have discovered this particular weakness, or that it will remain undiscovered if these academics are restrained from publication. Although the court briefly touches on (and rejects) the defence of public interest that would over-ride a prohibition against publication, the court did not expressly weigh in the balance of "gravity" the idea that a delay in publication might actually enhance the risk of future car theft. Each month that passes without publication is another month that secure system designers are unable to study whatever weakness has been discovered. In focusing its attention on the millions of cars that already use Megamos Crypto, the court seems to ignore the risk that millions of cars that are soon to be manufactured might otherwise be deprived of a more secure immobiliser system that is not being designed because cryptographers have been delayed access to this paper. (On this point alone, we suggest that even if a full trial on the merits produces a continuation of this injunction any such court order must surely be limited in time.)

In summary, the decision lacks a searching analysis of the degree to which this publication could be called "particularly grave". This lack of quantification of risk leaves academics (and publishers generally) with no real guidance on when publication can be properly enjoined especially where, as here, the court suggests that the claimant's case is otherwise in some doubt.

6.3 The source of the secret, and the burden of proof

Although the court's decision seems to hinge significantly on the legitimacy of the Tango Programmer product, the discussion of this product and its Bulgarian manufacturer is rather difficult to follow.

The court states clearly that it does "not know how Tango Programmer was created". (Para.34.) There is nothing in the decision to suggest that the manufacturer was present to answer questions about Tango Programmer. The opinion alludes to two pieces of evidence on this topic: the product description obtained from the Bulgarian company's web site, and certain observations made by the researchers in their paper.

Statements from the Bulgarian website seem to have prompted a great deal of mistrust by the court. It is not made clear why this was the case and the statements themselves (regrettably) are not repeated in the decision.

In their paper, the academics provide their opinion that the functionality of Tango Programmer goes beyond that which they believe is reasonably necessary for legitimate use. (The researchers made two additional points quoted by the court that are worth repeating. First, Tango Programmer is not the only tool they believe to be "overly" functional in this fashion. They cite a

tool called "AVDI" as another example. Second, the researchers note, "none of these tools is able to recover the secret key of a transponder or perform of (sic) crypto-analysis".) (Para.35, quoting from the unpublished academic paper.)

There is a sad truth concerning products that are offered for sale to locksmiths and security experts that seems to have been lost in the discussion. These products can be used to conduct activity that is perfectly legal, and they can be used in the commission of crimes. We suspect that this has been true for as long as the locksmith profession has existed.

This leads to what we suggest is a significant distraction in the analysis presented: conflating the issue of how Tango Programmer might be used with the issue of how the manufacturer of Tango programmer came into possession of the Megamos Crypto algorithm.

There is little doubt that Tango Programmer can (in the hands of a criminal) be used improperly. The court, however, seems substantially to ignore the idea that products like Tango Programmer have legitimate uses in the hands of locksmiths and automobile mechanics. This focus on how Tango Programmer can be used seems to have distracted from the enquiry at the heart of the decision: the method used to procure the algorithm embodied in the product. Even if the product is sometimes used by criminals, and even if the manufacturer is aware of this, we suggest this should not lead automatically to the conclusion that the manufacturer gained access to the algorithm through some form of industrial skulduggery.

There are many lawful avenues (other than licensed disclosure) by which the manufacturer of Tango Programmer could have obtained the algorithm. The court dismissed out of hand the possibility that the manufacturer paid for chip slicing. We find it difficult to understand how the court reached this conclusion with such assurance. Given the sale price of the Tango Programmer product at €1,000 per unit (and based on their web site, an additional periodic refresher fee to enable continued operation of the product) they may have decided that commissioning a one-time chip slicing project at €50,000 was economically viable.

There are other possibilities that the court does not address in its decision, and we are left to wonder if these were considered at all. Perhaps the manufacturer outsourced the chip slicing research to an overseas laboratory able to conduct this activity at a lower cost. Perhaps they did not pay for 100% of the reverse engineering effort. They might instead have paid a much smaller fee to a laboratory that had already reverse engineered the chip and then charged for access to such information. Or perhaps (like the academic researchers) they simply reverse engineered the software embodied in some other transponder programming product.

With little or no evidence on this issue we are left to speculate. And this is the point: the only evidence detailed in the decision leads only to speculation. And yet the court was persuaded that the manufacturer's apparent knowledge

about misuse of Tango Programmer by some was sufficient to infer – to the point that the court describes it as "obvious" – that the product resulted from trade secret misappropriation. (Para.36.) The court concludes its analysis of this issue with the statement that "it is probable that the claimants will succeed in showing that Tango Programmer's origin was not legitimate and that the [academics] ought to have appreciated that." (Para.39.) Without further published evidence on the point, we find the strength of this preliminary conclusion surprising.

As this is a significant point for persons engaged in research to consider, we note in passing that the court appears to accept that the burden of proof on this issue falls on the complaining parties. We infer this from the court's statement: "it is probable that **the claimants will succeed in showing** that Tango Programmer's origin was not legitimate…". (Para.39, our emphasis.) This suggests that it is first the responsibility of the persons attempting to enforce trade secret rights to prove that Tango Programmer was manufactured using a misappropriated trade secret. The court appears to have taken the view that the evidence on this point is sufficiently strong such that the academics then became responsible to produce evidence to the contrary.

6.4 The state of the academics' knowledge

The academics, of course, are not accused of misappropriating a trade secret from the creator or a licensee of that secret. It is instead suggested that they procured and reverse engineered a product (Tango Programmer) that resulted from some other person's misappropriation of a trade secret. The liability of persons who are accused of misusing confidential information that was misappropriated by another usually hinges on the state of their knowledge of that misappropriation. The court very clearly engaged with this issue. It ultimately concluded (on a preliminary basis) that, "it is probable that the claimants will succeed in showing that Tango Programmer's origin was not legitimate **and that the [academics] ought to have appreciated that**." (Para.39, our emphasis.) We believe the court's engagement with this issue would have benefited from additional explication.

The Supreme Court of the United Kingdom has recently discussed the liability of a third party accused of using information misappropriated by another. The Supreme Court's judgment in *Vestergaard Frandsen A/S et al v Bestnet Europe Ltd et al*, [2013] UKSC 31 (22 May 2013) was handed down only three weeks before the initial hearing in this case and five weeks before this decision was issued.

That case concerned a Mrs Sig, who had been employed by Vestergaard – a manufacturer of long life insecticidal nets. These nets were manufactured using certain techniques that were trade secrets. Mrs Sig resigned as an employee of Vestergaard and set up a competing business with two other persons who had been employed or engaged by the same company (one of whom was the scientist who effectively invented major portions of the trade secret). The trial court found that these two individuals misappropriated

Vestergaard's trade secret and used it to develop a competing product that the new business manufactured and offered for sale. These two persons were ultimately found liable for that misappropriation. The trial court found that Mrs Sig did not have actual knowledge of the misappropriation, but nonetheless held her liable for misuse of a trade secret belonging to her former employer. The Court of Appeals reversed on this issue.

The Supreme Court agreed with the Court of Appeals that Mrs Sig was not liable. In the Supreme Court's judgment, Lord Neuberger rejected a number of arguments advanced in favour of Mrs Sig's liability which finally included these: (i) that Mrs Sig had "blind-eye knowledge" of misappropriation, and (ii) that in setting up a competing business with two persons who had such close involvement with her former employers' trade secrets, her conduct suggested that she was "playing with fire" (Lord Neuberger's choice of words) and she should therefore expect to be liable. The Court rejected both of these arguments. The judgment makes clear that the liability of someone like Mrs Sig – accused of misusing a trade secret that had been misappropriated by another person – based on the theory that she turned a "blind eye" to misappropriation could only succeed if the complaining party could prove dishonesty on her part. The case referred to by the court with regard to proving dishonesty at this level involved shady business dealings by the constructive trustee of funds collected by a travel agent for an airline. (*Vestergaard*, para.42, citing *Royal Brunei Airlines Sdn Bhd v Tan* [1995] 2 AC 378.) The judgment further noted that her willingness to "play with fire", while it might assist in drawing an inference of dishonesty, was not enough on its own to fix Mrs Sig with liability. Lord Neuberger concluded, "if one plays with fire, one is more likely to be burnt, but it does not of itself mean that one is burnt." (*Vestergaard*, para.43.)

It seems to us that the *Vestergaard* decision has clarified the standard of what a complaining party must prove about the state of a third party's knowledge before that third party can be held liable for misuse of a trade secret misappropriated by somebody else. *Vestergaard* sets a very high bar on what must be proven. Actual knowledge of misappropriation clearly will suffice, as it always has. If it can be shown that the third party turned a blind eye to misappropriation – for which evidence of actual dishonesty is required – that will be enough. But anything less seems insufficient.

Returning to the current decision, the court made its decision on the basis that the complaining parties will probably be able to show: (i) that Tango Programmer was built with a misappropriated trade secret, and (ii) that the academics "ought to have appreciated" this. (Para.39.) The second prong of this conclusion may seriously understate what the complaining parties are, in fact, required to prove with respect to the state of the academics' knowledge. *Vestergaard* suggests that finding the academics liable would instead require the complaining parties to prove that the academics **actually knew** of the misappropriation, or that they **conducted their activities in such a dishonest fashion** that it demonstrates they turned a blind eye to the misappropriation. The court does not appear to have applied this higher standard in its analysis.

6.5 Responsible disclosure and the timing of this preliminary hearing

The decision constructs a narrative about the academics that is very unflattering. Faced with a request to delay publication for just a little while longer the researchers instead demanded the ability to publish immediately and thereby jeopardised the security of millions of cars. (We have already questioned whether this was such a serious risk.) While the court admits that the failure to make Volkswagen aware of the problem was not their fault, it chastises them anyway for failing to consent to any more delay: "A responsible approach would be to recognise the legitimacy of the interest in protecting the security of motor vehicles." (Para.41) The court delivers some of its most harsh commentary in describing the responsible disclosure process. "I think the defendants' mantra of 'responsible disclosure' is no such thing. It is a self-justification by defendants for the conduct they have already decided to undertake and it is not the action of responsible academics." (Para.42)

We suggest that a review of the evidence disclosed in the decision also supports a different narrative. This begins by considering the difficult work undertaken by the academics as part of their mission to support security research. The selection of Megamos Crypto as a potential research subject, the sourcing of Tango Programmer, the reverse engineering work needed to liberate the algorithm from the software, and then the core research work of examining the crypto algorithm for flaws. The decision does not state how long the researchers spent on this process, but we have little doubt that it was significant. Acting under ethical guidelines regularly applied by academics in this field, they approached the chip manufacturer EM with their findings in November 2012. They offered their assistance to develop work-arounds or replacement technology. They planned to publish their findings in August 2013, nine months after private disclosure. Having been open with EM, they heard very little in response. The researchers were then surprised when Volkswagen entered the picture in May 2013 – seven months after initial disclosure to EM – and sued them. Volkswagen requested and received an emergency temporary injunction with no notice to the academics. Given that the only meeting about the weakness in Megamos Crypto described by the court took place in June 2013 – a matter of weeks before scheduled publication – we are left to ponder how much emotion may have entered the situation at this stage.

The difference in these two competing narratives demonstrates a significant disagreement about what constitutes "responsible disclosure". It appears that the court may not have fully appreciated how this phrase is used as a term of art in the context of security research.

There are three main methods of public disclosure that are in common use in this admittedly abstruse field: (1) non-disclosure, (2) responsible disclosure and (3) full disclosure. In the first case, the researcher tells the affected party, and then says nothing more; in the third case, the researcher publishes without telling the affected party in advance and without any regard to their

interests. The second way is a middle path between these extremes that is now very widely followed by academics and more generally security researchers. Typically, six weeks is set as the "time to disclosure" in the case of software flaws, and six months in the case of hardware flaws. However, in extreme cases, where no simple fix is available and the impact is very serious, researchers might feel compelled to wait longer than six months.

These time scales (six weeks and six months) are not unique to these academic researchers. They are widely used baselines within the field of security research. We imagine that the researchers felt that they had already "gone the extra mile" by disclosing nine months in advance of publication, and might have felt rather abused when someone other than the product's manufacturer suddenly appeared and brought a lawsuit only two months before planned publication.

It is crucial to understand that "responsible disclosure" is simply a phrase used by researchers to describe one approach to the public disclosure of security flaws, one that is certainly more responsible than full disclosure, and arguably even more responsible than non-disclosure, given that the latter approach does not create any incentive for the affected party to address any disclosed flaws in their products. The court did not appear to appreciate this distinction, given the way in which the decision criticizes the researchers. (Para.42.)

Furthermore, and more importantly, it is apparent (from para.14) that the court's understanding of the term is incomplete: there, a definition of responsible disclosure is offered which entails "telling the manufacturer in advance" about the flaws, but which does not include the critical point that, in this mode of disclosure, a date is set up-front for when disclosure **will take place, irrespective of the circumstances at the time when that date is reached**. Establishing such a publication deadline when disclosing to the manufacturer is not simply the arbitrary or capricious act of a petulant researcher. This mechanism is used to prevent affected parties (who, as noted above, often form part of complex supply chains) from unnecessary dithering and to ensure they have an incentive to address the identified security flaws. It seems that this missing point concerning timing is what leads the court to heap opprobrium on the researchers in paras. 41 and 42, where it is opined that "it was not consistent with the idea of responsible disclosure for the defendants to simply say, 'We are going ahead anyway'." and "I think the defendants' mantra of 'responsible disclosure' is no such thing." There is a value judgment implied by the use of the word "mantra" here – this meaning a phrase repeated often and without significant thought. Our experience is that academic security researchers and industrial consumers of cryptography alike do understand the significance and methodology of responsible disclosure, and accept it as the preferred, if not universal, *modus operandus* for disclosing security vulnerabilities. This apparent breakdown in understanding seems to heavily colour the court's view of the academics' probity.

We find the strong language used to describe the actions of the academics both puzzling and disappointing. First, it is clear that their approach to

"responsible disclosure" was well within normal guidelines followed by security researchers for the benefit of the security industry (and society) as a whole. Even if it were not, the strength of the court's condemnation is surprising given the reality it had already acknowledged – reasonably accessible methods are available that would allow the academics or anyone else to publish the algorithm without the permission of the complaining parties.

6.6 Delay in enforcement

As noted above the academics are not themselves accused of misappropriating a trade secret. This decision is based on the theory that they have (through reverse engineering) obtained access to a secret that was misappropriated and then incorporated into a product by some other person more than four years before this case was filed.

This raises an issue that was not explored in the decision – the rather long period of time it appears to have taken the complaining parties to assert their rights in the allegedly confidential algorithm.

If (as the court believes to be highly probable) Tango Programmer was built using a misappropriated trade secret, this misappropriation took place sometime before 2009. The decision does not suggest that the Tango Programmer product is sold secretly behind closed doors. The court, for example, relies heavily upon statements made by the manufacturer on its web site.

We reviewed the product web site (www.tangoprogrammer.com) and discovered that it includes address details and phone numbers of resellers around the European Union. It also makes clear that the product can be used to analyse and code transponders that employ various versions of the Megamos Crypto system. (We also reviewed relevant pages in the Internet Archive "Wayback Machine", which confirms that these details were also posted at this site as early as May 2012.)

Given that the Megamos system is used in many millions of automobiles, and given the view put forward that Tango Programmer can be used by criminals (among others), we are left to ponder why the decision makes no mention of the claimants filing any prior enforcement actions against the manufacturer of Tango Programmer (or any other third party device that might embody the algorithm without license). The law does not only give the complaining parties the right to forbid publication of the secret – they could also (for example) demand an accounting of profits from companies who use a misappropriated trade secret for commercial gain, or indeed obtain on order prohibiting further sales of infringing products. Given the centrality to this case of the factual question about the provenance of the algorithm in Tango Programmer, if such prior actions had been filed we assume that the current decision would have made some reference to them. (Indeed, since the manufacturer openly advertises a UK office address one assumes that they might have been

named as an additional defendant in this case) As it is, we are left to infer that no such prior enforcement action has occurred.

Any such enforcement delay could be legally significant. First, courts expect parties to pursue their rights with reasonable speed and not to delay their claims unnecessarily. Secondly, in balancing the equitable rights of the parties in a case like this the court is within its rights to enquire about the reason for delay and the motivation of the complaining parties in pursing action now – against the authors of a paper the contents of which might later be used to facilitate criminal activity, rather than against the manufacturer of a device that they complain (and the court accepts) is already being used by some to facilitate criminal activity.

This point about delayed enforcement does not appear to have been raised with, or considered by, the court.

7. Impact on academia and public policy in the UK and beyond

This ruling is likely to have a chilling effect on legitimate security research in the UK. While the circumstances of this case are rather specific, and the decision hangs on those specifics, the case creates a degree of uncertainty and confusion around what can, and cannot, be done by security researchers without running the risk of encountering legal obstacles. For academic researchers, "publish or perish" is a no less pressing or relevant a motto for it being hackneyed through overuse. And the investment in time and effort required to conduct the kind of research relevant to this case is significant, as are the risks that any given research avenue selected will turn out to be unfruitful. So the mere perception that legal barriers to publication *might* arise is likely to cause some researchers, particularly new entrants to the field, to think twice about starting at all.

It is then especially ironic that, all the while, the UK government, through its funding agencies (such as EPSRC) and UK government departments (such as CESG/GCHQ and Business, Innovation and Skills, BIS), has been investing heavily in cyber security research, with a proportion of that funding being directed towards projects involving the development of techniques for the analysis, discovery, and eventual elimination, of weaknesses in security systems.

We may also speculate that the ruling may have repercussions beyond the UK. Academic research in cryptography and security is a discipline now observed routinely around the world. Multi-country collaborations (like the collaboration that is the subject of this case) are commonplace. It is unclear whether the High Court of England would have been vested with jurisdiction of this case but for the fact that one of the authors and his employer are resident in the United Kingdom. The remaining two authors are normally resident in the Netherlands. The putative publisher is based in the United States. Certainly courts in the United States are highly suspicious of such prior restraint cases due to a combination of the guarantee of free speech (under the First Amendment of the US Constitution) and certain limitations in the US treatment

of trade secrets. (See generally, Samuelson, "Principles For Resolving Conflicts Between Trade Secrets And The First Amendment", 58 Hastings L.J. 777 (March, 2007).)

As a result of this decision, it seems plausible that researchers based outside the UK may be less enticed by the prospect of working with UK-based researchers given the possible injunction of their eventual joint research papers. The effect would be to isolate UK-based security researchers, at a time when national governments are strongly emphasising the need for cross-border efforts in cyber security research (see for example the UK Cyber Security Strategy at: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/60961/uk-cyber-security-strategy-final.pdf).

8. Conclusion

In granting a preliminary injunction that partially restrains publication of academic research into weaknesses in the Megamos Crypto system, the English High Court has taken a step that is – and should be – troubling to legitimate security researchers.

In our opinion, the court's decision evinces a lack of understanding of the foundational principles of cryptography and secure system design that would have been necessary to conduct an appropriate enquiry into the risks of publication. The decision also appears to lack a clear understanding of the term of art "responsible disclosure", and the well-established role that this plays in security research.

Although this is a preliminary decision, given the admitted infringement of free speech we find the application of law to the facts in this decision to be surprisingly brief and unhelpful. We are especially puzzled by the court's willingness to jump so quickly to the conclusion that the manufacturer of the Tango Programmer product engaged in misappropriation of a trade secret, and having reached that conclusion that the academics ought to have been aware of the misappropriation. If the court had better reasons to draw these inferences from the preliminary evidentiary record, it is unfortunate that the court did not describe this evidence in the published decision.

We are also troubled at the chilling effect that this decision may have on legitimate security research in the UK. This decision, which we expect will be viewed as out of step with the prevailing trends of other countries regularly engaged in such research, could have the effect of isolating UK security research academics from their international colleagues – at precisely the time that the government in the UK is encouraging an increase in such research and in international cooperation.

As a final comment, we have no doubt that the judge in this matter – who was required to hear this application and make this decision in a very compressed time frame – is an extremely able jurist. Judges, no matter how able, cannot be experts in all subjects. In English courts (and other common law courts

around the world) it is the responsibility of others to explain to the court key elements of technology under review. Perhaps for no reason other than the compressed timetable leading up to the hearing and decision, it appears to us that this process of explaining complex technical facts and practices from an otherwise abstruse specialist field has somehow broken down.

About the Authors:

Robert Carolina is a Director with the Origin law firm in London, and a Senior Visiting Fellow with the Information Security Group at Royal Holloway, University of London. He is a graduate of the University of Dayton (BA Political Science, 1988), Georgetown University Law Center (Juris Doctor, 1991), and the London School of Economics and Political Science (LL.M in International Business Law, 1993). He is admitted to practice as a Solicitor in England & Wales, and as an attorney-at-law before the US Supreme Court and the courts of the US State of Illinois. In his legal practice, Robert advises clients of all sizes with respect to creating, protecting, procuring, licensing, selling, distributing, deploying, using, and managing a wide variety of information and communication technologies – including cryptography-based products.

Kenneth G. Paterson is a Professor of Information Security and EPSRC Leadership Fellow at Royal Holloway, University of London. He holds a Ph.D. in Mathematics from the University of London. He is a member of the editorial board of the *Journal of Cryptology* and was Programme Chair for Eurocrypt 2011. He publishes widely in the fields of cryptography and secure protocol analysis, is in frequent demand as an invited speaker, and consults widely.