Edinburgh Research Explorer

A Framework for Modelling Trojans and Computer Virus Infection

Citation for published version:
Thimbleby, H, Anderson, S & Cairns, P 1998, 'A Framework for Modelling Trojans and Computer Virus Infection', *The Computer Journal*, vol. 41, no. 7, pp. 444-458. https://doi.org/10.1093/comjnl/41.7.444

Digital Object Identifier (DOI):
10.1093/comjnl/41.7.444

Link:
[Link to publication record in Edinburgh Research Explorer](#)

Document Version:
Publisher's PDF, also known as Version of record

Published In:
The Computer Journal

OPEN ACCESS

A Framework for Modelling Trojans and Computer Virus Infection

Harold Thimbleby[1], Stuart Anderson[2] and Paul Cairns[1]

[1]*School of Computing Science, Middlesex University, Bounds Green Road, London N11 2NQ, UK*
[2]*Department of Computer Science, Edinburgh University, King's Road, Edinburgh EH9 3JZ, UK*
Email: harold@mdx.ac.uk, soa@lfcs.ed.ac.uk & p.cairns@mdx.ac.uk

It is not possible to view a computer operating in the real world, including the possibility of Trojan horse programs and computer viruses, as simply a finite realisation of a Turing machine. We consider the actions of Trojan horses and viruses in real computer systems and suggest a minimal framework for an adequate formal understanding of the phenomena. Some conventional approaches, including biological metaphors, are shown to be inadequate; some suggestions are made towards constructing virally-resistant systems.

1. INTRODUCTION

Computer viruses are not merely an irritating and destructive feature of personal computing, they also mimic the behaviour of biological pests. Computer viruses hijack the facilities provided by the host computer, lie dormant, cross-infect other hosts and when activated cause various forms of damage, from obvious destruction of host data to more subtle changes that are much harder to detect.

The predominantly human medical metaphors employed when discussing computer viruses are misleading because of ethical connotations. The word 'virus' itself is a rich metaphor and terms like 'infection', 'viral attack', 'disinfectant' and 'viral damage' are frequently used of computer viruses. While these terms lack any precise meaning in relation to computers, we know roughly what they mean. Furthermore, the debate about the computer virus threat is predominantly experiential: awareness of viruses is biased towards frequent infections caused by successful viruses; debate centres around newly identified virus cases.

In order to begin to tackle viruses effectively we should define the essential, general features exhibited by viruses without appealing to metaphor or limited experience to fill in the gaps in our understanding. Computer viruses seem to be a simple phenomenon (hackers manage to write them in quantity and there are virus construction kits that anyone can use to develop sophisticated viruses) yet they do pose a real hazard. The science fiction anticipation of viruses was thought to be obvious [1]. It is of interest, then, to explore the fundamental limits of detection and prevention and other issues. On the face of it, one would look to standard models of computation, such as Turing machines, as a starting point for this investigation. However, a closer analysis requires concepts not usually considered in the standard models of computation. We agree with Wegner

[2] that Turing equivalent formalisms are not sufficiently expressive for systems that interact. A virus has to *enter* a system to infect it and this is such an interaction. And as we shall argue, the biological metaphor is inadequate too (though it can still inspire computational ideas, both good and bad). The computational process undertaken by a virus does not terminate; moreover, when the (possibly) Turing process of an infected program terminates, the virus has already infected another program that will subsequently execute—a virus that terminated would not be infective.

Viruses should not be dismissed as a trivial problem for users of computers that do not run secure operating systems. Together with Trojan horses, they are a major menace for any sort of system and are a particular concern for computer systems connected to networks [3]. As well as being a practical problem, this paper will show that viruses also pose theoretical problems. Indeed, this paper, by laying out some of these problems, begs many questions that raise many further research questions. Some of these research questions will be pointed out explicitly throughout the paper.

In summary, this paper highlights the limitations of applying the conventional computational models to computer virus phenomena and it offers a new framework. We motivate our paper by defining and naming the phenomena under discussion. We then examine a variety of related phenomena, such as: compiler viruses, Cohen's proof about virus detection, the possible structures of virally resistant systems, and finally, Koch's Postulates and other biological issues. The reason why people write viruses is beyond the scope of the present paper; see [4] for an introduction and overview of the literature.

1.1. Orientation and terminology

Viruses and Trojan horses make computers do things that their users do not want them to do. The term Trojan horse

is due to the Greeks' use of a hollow wooden horse filled with warriors to deceive the Trojans. The Trojans were violently torn between accepting the horse or rejecting it; indeed, they did start checking it, but were fatefully tricked into tearing down their own city walls and pushed it in themselves [5]. More recent discussions of viruses include [6, 7, 8, 9, 10, 11]. The interested reader is referred to [12] for up to date information on *specific* viruses, which are not the concern of the present paper.

For clarity and to avoid an air of pedantry we abbreviate 'Trojan horse' by the single word 'trojan' when referring to computer programs. Trojan programs, sometimes with names like `sex`, often seem so attractive that users are sooner tempted to 'bring them into their city', like the original Trojan horse, than to test or to reject them—so bringing upon themselves the results intended by the trojans' designers. The complication is that trojans do something unwanted yet they offer to provide a wanted service. One might start by considering four simple categories of trojan:

Direct masquerades pretend to be normal programs.

Example: a program called `dir` that does not list a directory, the normal use of the command of that name. Computer systems often permit many programs of the same name to co-exist, with consequent opportunities for masquerading trojans.

Simple masquerades do not masquerade as existing programs, but rather masquerade as possible programs that are other than they are.

Example: programs with names like `sex` above.

Slip masquerades have names approximating legitimate program names.

Example: a program called `dr` that might be activated if the user mis-typed `dir`. Since users want to install programs (e.g., for maintenance or upgrade purposes) and perhaps write their own programs, in practice few systems provide restrictions on the names of programs; even if they do, there must be fewer restrictions on the programmers who design them. The consequence of this liberality is undetectable/unidentifiable trojans.

Environmental masquerades are not easily identifiable programs invoked by the user, but are typically already-running programs that provide an unwanted interpretation of the user's commands or other activities.

Example: an operating system whose login prompt to the user is an otherwise clear screen and the prompt 'login:' can be trojanized by constructing a program that intercepts the user, by clearing the screen and issuing the login prompt itself. It can then embark on any activity it pleases, typically recording the user's name and password; to avoid detection by the user it would then transfer control to the authentic login program. (If this is not feasible, it could misleadingly report a password error, ask the user to try again and then terminate so that the authentic login program takes over completely.)

Example: when an entertainment CD is inserted, it may start playing automatically—if it contains executable code, it may be executed and do (or prepare to do) damage. The Macintosh `AutoStart 9805` is a trojan of this sort; it commences execution as a side-effect of inserting removable media.

Beyond these basic categories, it is fruitless providing precise definitions of the categories or types of unwanted behaviour of trojans (or, indeed, viruses) since their (i.e. their programmers') intention is to be devious. An actual trojan may exhibit mixtures of these behaviours and random effects to confuse and mislead the user. There is one trojan which, itself doing no direct damage, instructs the user to switch off the computer (pretending to have detected a serious hardware fault); if the user responds to this advice, they are liable to lose any data that has not been saved to disc. If we admit these programs that merely *say* things, what should we make of programs constructed by experimenters that say what they *would* do had they been actual viruses or trojans?

Given the wide range of behaviour, both intended and accidental, not surprisingly the terms trojan and virus are the subject of lively debate [13], based on a number of architectural and other distinctions. Most work in the area has concentrated on purely representational concerns, where every 'variant' of a virus technique (such as the Vienna virus represents) are considered different. (This has resulted in particular interest in 'polymorphic' viruses, which modify their representation on each infection.) Our approach is more semantical. Whatever its behaviour, a viral infection has three core components:

A trojan component. An infected program does something unwanted in certain circumstances. The trojan component is sometimes called the *payload*.

A dormancy component. The viral infection may conceal itself indefinitely. Trojans, too, may use dormancy to conceal their presence, but with a virus dormancy (or, equivalently, unnoticed trojan damage) is essential for the effectiveness of their third component.

An infective component. Infected programs infect further programs, which then behave in a similar way. (Viruses may wish to avoid re-infection, because re-infection takes time or space and may therefore make an infection more obvious. Viruses often include a heuristic for self-detection, a procedure which, if identified, might be directed against them.)

Even this abstract list has problems! A virus that only replicates consumes computer resources (it uses disc space, it wastes time . . .) even though it may not have an explicit trojan component. Or a virus that is never dormant might still be able to infect, provided it uses some other means of avoiding detection.

If we allow programs to run on various machines, across platforms, why stop at digital computers? Are viruses with no infective component, but which trick humans into

spreading them, rather than spreading themselves, really viruses? Are chain-mail, spam [14] and email panic scares, which are replicated as people email the panic to people they wish to warn, really viruses [15, 16]? Certainly, these are serious problems. As communications systems become more sophisticated, performing actions (such as forwarding email) on behalf of users, the traditional distinctions between trojans and viruses become harder to maintain.

In particular, the trojan component may be 'harmless' (viewed inside a computer) but precipitate a chain of behaviour *outside* any computer, ending with a harmful conclusion. Although the trojan component may run at a definite time, the user may in principle be unaware of the consequences until much later. A particularly pernicious program, ADTrojan, sends rude email to other users. The only action a computer could detect would be the (not necessarily) unauthorised mailing. The recipients of the email, however, can complain, leading to repercussions for the original user. Other sorts of delayed trojan activity include programs that steal passwords: the noticeable damage occurs much later when the passwords are used to enable third parties to do other work—which of course may be part of a larger chain of activities, only the end of which is blatant damage. The legal consequences of this sort of indirect behaviour involving innocent parties have not yet been worked through satisfactorily.

We might also add that a virus has a **survival component**. Crucial to the long-term practical success of a virus is that it can infect in some sense faster than it can be eliminated; to do so, it may infect other computers—these other computers are often the responsibility of other people or organizations, and so the elimination of a virus requires human co-operation, which is typically slower to recruit than the spread of the virus. Clearly, virus survival is of practical significance, and the organizational structures anti-virus manufacturers use to distribute their products is of commercial interest. However, survival is an emergent property, that follows from dormancy and infection, so we do not formalize it here. We consider it an epidemiological issue (see, e.g., [17]).

Because we use words like virus and trojan, and often give them names (e.g., *the* Vienna virus), it is easy to think the problem is just one of identification.[1] However, bugs, viruses and worms can, under appropriate assumptions, be benign. Bugs are used constructively in certain forms of AI programming; viruses can be used to spread information [19]; worms can be used to distribute processing [20]. Moreover, the people who write destructive programs may think their programs are benign, even if others disagree. Clearly identifying 'the' problem is a human issue, related to what one wishes to define as a problem. From the formal point of view, the question is, if one wished to classify something as a problem, whether that thing could

be distinguished to be so identified. Of two programs, before one can say one is a trojan, it has to be established that they are different.

Thus to develop our framework, we wish to be able to model the essential issues, but perhaps not make all of the distinctions some authors have made. An adequate framework must involve notions of behaviour, invisibility, infection—what is the 'same' virus? These questions have practical implications: How should viruses be named? Are all strains of the Vienna virus the same virus, or are they different?[2]

A framework must address, or relate to, the confusion that can exist over the interpretation of 'damage.' If a user, being malicious, *intends* damage, are we then to consider the activities wrought by a trojan on his or her behalf as 'constructive'? The notion of damage clearly assumes particular value systems and intentions on the behalf of particular classes of user. Indeed, the behaviour of a user developing a compiler system, who in the course of their legitimate work, compiles code, replaces system resources and recompiles various programs and so forth, is hard to distinguish on purely technical grounds from a user who replaces programs with malicious intent. (A tongue-in-cheek article argues that research ideas behave like viruses [22]: which explains the proliferation of research publications and journals.)

Although both trojans and viruses may be difficult to isolate after damage has been detected, the additional components of viruses ensure they are not so easily eliminated. As a result of dormancy and infection, a virus normally makes many copies of itself (not necessarily identical) before the damage is detected. The virus may make copies of itself essentially outside the computer it is running on: it may make copies of itself on removable media (such as a floppy disc) so that the user, on putting the disc in another computer, infects a remote computer; or it may use networks to do the infection more directly itself. When—if—the virus is detected it will already have spread somewhere else, where its management is someone else's responsibility.[3]

We assert that a virus can only be defined semantically, that is in terms of the meaning of programs, rather than in syntactic patterns appearing in code. This belief is strongly at variance with most of the literature on the topic [23, 24]. All viruses to date have been trivial, and this gives the impression that viruses might be classified as syntactical, with various identifiable structures—for example, fragments of code that correspond with the three conceptual components listed above. With the prevalence of so-called polymorphic viruses (ones that have different syntactical forms [25]), this view is of course questionable, but there is a lingering (and incorrect) view, perhaps partly

---

[1]This is what most virus detection programs do: they look for viruses whose signature they recognize. If a user (or a program they are running) knows what virus to look for, there is a partial oracle for (what we will show to be) the non-computable function that identifies the virus. See [18].

[2]The Vienna virus has many variants, partly due to the publication of source code for one variation of it in [21].

[3]It follows that successful elimination of viruses requires a distribution mechanism for the antidotes: this is usually done by conventional marketing channels (that do not use the same vectors as the virus), rather than by virus-like computer-based replication.

inspired by the relatively concrete notion of biological chromosomes or genes, that there 'really' is a syntactic basis. If any biological metaphor is required, computer viruses are more like Dawkins' 'memes', viruses of the mind, than biological genes [26], a point we return to later (Section 5). In our view, what is important about a virus is not how it works, but what it accomplishes.

A virus is best defined by saying what it is to be infected. An infected version of a program p is a program p' that behaves indistinguishably from p on 'most' inputs, behaves distinguishably from p on some inputs and sometimes when p' is executed it infects other programs in the 'same' way. This preliminary definition begs a lot of questions and uses terms without definition. The remainder of this paper, then, is directed to making these notions precise. However, one thing is clear from this definition, the notion of viral infection is not linked to any particular syntactic representation as code.

Our paper will necessarily introduce a new definition of virus, which we will provide after justifying a new framework within which to express the definition. In particular, so-called worms, which some authors distinguish from viruses as autonomous programs rather than programs physically carried by users, are not distinguished in our framework. We will find it convenient to use the general term *infection* to indicate the presence of a trojan or a virus: an infection is, from the hacker's point of view, an intentional bug. In keeping with our abstract approach (and indeed the medical usage) this does not imply a particular manner of acquiring the infection, nor that the infection is transmissible to other programs or computer systems.

1.2. Previous work

We are aware of previous work on the theoretical limits to the detectability of trojans (e.g., [27, 23, 28]). Our own first undetectability proof for masquerades [29], based on this work, was questioned for not explicitly addressing security issues [24] (a criticism that also applies to the other methods). We have also attempted to identify and distinguish the causes of co-operative behaviour and destructive behaviour [30], but this was informal; and we have found problems in a proposed detection mechanism [31], but our findings did not constitute a general theory.

2. INADEQUACY OF TURING MACHINE MODELS

It might seem obvious that personal computers are Turing machine equivalent and, as personal computers get viruses, so Turing machine models are an appropriate basis for studying issues of infection. However, this view is false.

First, we will put forward a plausible argument, then we shall provide a proof. Turing machines are infinite machines, whereas personal computers are finite. Clearly, Turing machines are an idealization. Personal computers have properties not shared with Turing machines. Some real-world properties of personal computers—plausibly including the issues of trojans and viruses—are not modelled by the idealization.

Now, more formally, suppose we define 'infected' as a predicate on Turing machine programs and we do not try to formalize what can be observed as a result of the execution of the program. What can this mean? It can only mean that the result of running the program is either unchanged (the infection has no observable effect) or that the result is incorrect (the infection is not 'hidden'). In neither case do we have a non-trivial problem. All previous work is based on abstract models of computation together with concrete notions of replication. Such approaches cannot capture various subtle notions of, for example,

- **masquerading**, where a user knows names of programs and anticipates their likely behaviour;
- **infection**, where viruses may encrypt themselves and be different each time they infect.

Our comments will apply equally to any Turing equivalent model of computation, such as λ-calculus. Note that it is important to distinguish between a representation of viruses (clearly, any computable theory of viruses is directly representable on Turing machines) and a definition of an effective model.

There is a tension that, on the one hand, the conventional undetectability results are too pessimistic; on the other hand, naïve notions of infection may lead to excessive optimism—in that specific remedies may be attempted that rely on these notions of infection. Questions remain whether virus activity, replication and infection, can be usefully constrained in suitable computational models where such activity is explicitly but generally modelled. We will explore some of these issues in Section 4.

2.1. 'Other' programs

The notion of a virus infecting other programs has to be modelled: such a notion is essential to any discussion of viruses. However, 'other' programs cannot be modelled by a conventional Turing machine model since there is no other program—even if the Turing machine tape contains several interleaved programs whose execution is multiplexed.

2.2. Self-awareness of infection

In principle, a program could be infected by a virus in such a way that it could not tell it was infected (the malicious code could interfere with any infection testing): thus, the reliable judgement whether a program is infected depends on an external mechanism that is not affected by that infection. In fact, so-called armoured viruses [32] exist for the Intel 80486 processor, which detect attempts to single-step (and hence understand their behaviour) them by rewriting already-pipelined instructions so that they can distinguish whether they are being run directly by hardware or being interpreted.

We will avoid this problem by introducing the concept of trojan and viral methods, which are *outside the system*—they are outside the standard Turing model—whose infection is being considered.

2.3. Self-replication as a special case

If infected programs were empty, computer viruses would involve self-replication. Self-replication of programs has attracted a great deal of attention, particularly because of the apparently paradoxical problem of encoding concrete representations of programs *inside* themselves: that is, when a self-replicating program is run, its output is its own source code. From this point of view of self-replication, a program that merely accesses its representation (e.g., from a source code file) would be cheating!

Turing completeness (sufficiency for implementing a universal Turing machine) is neither necessary nor sufficient for self-replication.

It is not necessary since self-replicating programs can be constructed as non-recursive straight-line programs. A programming language could therefore be constructed that was not Turing complete, but which was sufficient to implement non-cheating self-replicating programs.

It is not sufficient, since Turing completeness is up to representation. We can construct a programming language that is Turing complete, but where the language uses symbols that running programs cannot output. We then have Turing completeness where self-replication is impossible. A proof of insufficiency where the domain and range of a Turing complete language are identical is given in [33].

Fokkinga [34] gives a construction for self-replicating programs and adds a variation: a program that recognizes its own source code—which is a prerequisite for a virus not to repeatedly re-infect a file (see Section 2.2). Kanada gives an example of a self-replicating program that runs on the World Wide Web [35].

2.4. Time and space

Turing machines are infinite machines and their speed of operation is immaterial. The real world of practical computers is finite; space is finite and time is significant. Since, for small n, every program on an n word (RAM + backing store) computer can be enumerated and classified as acceptable or infected by inspection, it follows that an adequate framework should allow for the complexity of classification.

Also, some viruses do their damage 'merely' by consuming resources. This reduces the effective speed of the computer or loses free space. Neither of these effects are of any consequence in a computational model that admits an infinitely fast, infinite memory machine.

3. A NEW FRAMEWORK

In any rigorous framework for the study of viruses there must be a mechanism to create and distinguish between various programs. Without loss of generality, the necessary extension is an environment, a mapping of names to programs, equipped with the appropriate operators to make enquiries of it and to manipulate it.

The object of our study is computer programs together with their inputs and outputs.

We could imagine a real computer to be an array of bits, including its RAM, screens, backing store and the state of its CPU. Parts of this array correspond to programs, parts to data files and various parts to such components as boot sequences and directory structures. The meaning of certain sequences of bits may depend on where they reside in the computer. A bit pattern may be a graphic image, but somewhere else it might be a program that can be run. After an error, such as stack overflow and obtaining a bad return address, the computer might be directed to interpret a graphic as a program.

We call all these things together (the full state of a machine, program texts, graphic images and so forth) the *representation*. The collection of all possible representations we denote R. Any given representation, $r \in R$, is finite. Note that R includes all fault conditions, such as the computer 'locking up'.

The user is not concerned with representations. Indeed much of the memory of a computer is hidden from the user, in boot sectors and other places. The user is concerned with names of programs and the computer uses those names, following various rules, to obtain the representations of the corresponding programs. We therefore introduce the *environment map*, E, which is a fixed map from a representation (typically the current configuration of the computer) to an environment, a name-to-representation map, which takes names (chosen from a countable set of labels L) and, if they are defined, obtains their corresponding programs or other data: $E: R \to (L \twoheadrightarrow R)$. The domain of $E(r)$, $\text{names } r = \text{dom } E(r)$, is finite and computable and (as is made clear later) will normally include some fixed names independent of r.

Note that 'names' are a very general concept and include, for instance, locations on a graphical user interface screen, or textual names in a conventional file store. In practice, the environment will have a structure that may have security implications, but this is not required for our framework.

Programs may be run and running programs usually changes the state of the computer. We say that the *meaning* of a program is what it does when it is run. If $r \in R$ is a representation that includes a program p, then $[\![p]\!]$ is its meaning: $[\![\cdot]\!]: R \to (R \to R)$. The meaning of a program is that, when run, it transforms representations into representations. Note that our approach admits everyday complexities such as operating systems, virtual machines, spreadsheet macro programs, dynamically loaded Java applets and so forth—but it is not necessary to model communication, non-determinism or concurrency to capture what viruses do. A more thorough definition of programs and representations could certainly be developed (and would be a useful research project), but for our purposes we do not think it would provide any additional clarity to do so—as more structure is introduced, it is very hard to avoid implementation bias and the associated obscurities of 'real' machines.

Where no ambiguity arises, we abbreviate the structure $\Omega = [R; E; [\![\cdot]\!]]$ by R.

Our framework does not require a virus to reside in 'a' program; conceivably it could reside in two or more co-operating fragments, none of which alone behaves like a virus. A trivial way to do this is by threading the virus code around the representation, but more subtle techniques are possible: see Section 4.4.

Crucial to a framework for examining infection is that programs appear, perhaps for long periods of time, to be other than what they really are. A user may find two programs to be indistinguishable even though they are not equal. We define these terms precisely as follows:

Two programs p and p' are *equal* when

$$\forall r \in R: [\![p]\!]r = [\![p']\!]r.$$

However, unlike identity, equality is not a computable relation,[4] and even to check equality of program outputs for a small collection of inputs would require the exhaustive examination of the entire state of the computer. More practically, in an attempt to take account of the amount of time one is prepared to devote to checking an output is what is required, we define *similarity* to be a poly log computable relation on R (see below), denoted $\sim$.

We do not assume similarity is an equivalence relation. In particular, similarity is not transitive: we may know $a \sim b$ and $b \sim c$ but also that $a \nsim c$, given the poly log time restriction. Since $\sim$ is computable it must be that either $a \sim c$ or that $a \nsim c$, and that this result is known in finite time: the computation of similarity may make 'mistakes'. There then arises the possibility that unequal programs are similar: although two programs are different, we may not be able to decide that in the time available. The key point is that similarity ($\sim$) of programs is not equality ($=$) of programs, for if it was there would be no serious problem in detecting different programs.

We define *poly log*, and a related quantifier, *for most*.

Poly log computable. Poly log computable is a restriction that a function can be computed in less than linear time on the total size of its arguments. Poly log is a requirement that a function of representations can be evaluated without examining the entire computer representation (which can be done in linear time). If the entire computer representation could be examined at every step in a process a number of detection questions become trivial; furthermore, it would implausibly suggest the user is aware of the entire configuration of the computer, including boot sectors, operating systems and so forth.[5]

$\mathcal{M}$ **(for most).** We need to introduce a new quantifier, *for most*, written $\mathcal{M}$. A technical definition of this notion is not required in what follows; a definition that captures the intuition and the relation with poly log computable (or some other measure) is a research project.

---

[4] The program fragments $1 + 1$ and 2, suitably interpreted, are equal but not identical.

[5] For small computers, say hand held calculators, the poly log restriction may make it feasible to examine the entire representation space.

Two programs are indistinguishable when they produce similar results for most inputs. If p and p' are two program representations they are *indistinguishable*, written $p \approx p'$, if and only if

$$\mathcal{M} r \in R: [\![p]\!]r \sim [\![p']\!]r.$$

We need some convenient notation.

- $\widehat{p}$. A representation or program $\widehat{p}$ is an attempt to trojan p; $\widehat{p}$ is to be taken as a metaname, and may, in fact, have no relation to p.

- $r \xrightarrow{l} r'$. We write $r \xrightarrow{l} r'$, iff $l \in \mathsf{names}\, r$ and $r' = [\![E(r)l]\!]r$; this extends naturally to finite sequences of program names $l_1 l_2 \ldots l_n \in L^*$.

- s/c. Let s/c be the object code corresponding to s when compiled by a compiler c. This seems intuitive enough, but it assumes that out of the entire machine representation it is possible both to choose the source code and the object code resulting from compiling the source. In fact, just before a compiler (or any other program) is run some other program (e.g., the operating system) places the parameter of the compiler in a known part of the representation; some convention in the program (no doubt enforced when it was compiled!) then specifies the location in the representation of its parameter. If we label these locations $\Lambda_1, \Lambda_2, \ldots$ then we have (allow s/c to be empty in the case that s is not well formed with respect to c):

$$\exists \Lambda_1, \Lambda_2 \in L: \forall r \in R:$$
$$s/c = E[\![c]\!]r\Lambda_2$$
$$\textbf{where } s = E(r)\Lambda_1$$

If c_s is a compiler c in source form, $c = c_s/c$. The notation extends naturally to finite sequences of applications of a compiler: $s_n/s_{n-1}/ \ldots s_0/c$.

With these preliminaries, we may now define trojan and virus. In attempting to do this we find that we cannot eliminate the environment from the definition. The notions of trojan and virus can only be understood relative to their binding environment. Hence, rather than define trojan and virus as such, we will define a recursively enumerable relation to capture the *method* ('infection') employed by the trojan or virus, respectively.

3.1. Trojans

Trojans may corrupt something unnamed (say, a boot sector) which when run at a later time results in an 'obvious' trojan effect—but even that 'obvious' trojan effect cannot usually be determined except by running a program, for example to check that certain files are still present.

As a first approximation, we might say that programs $p, \widehat{p}$ would stand in the relation trojan when there is some representation r that distinguishes their behaviour; informally, p trojan $\widehat{p} \Leftrightarrow \exists r \in R: [\![p]\!]r \nsim [\![\widehat{p}]\!]r$.

Notice that the trojan relation is symmetric: without assuming what p or $\widehat{p}$ is supposed to do, we cannot know

which program is intended as a trojan of which. We could imagine a hacker making the ironic comment that a real login program trojanized their subversive login program. Since it is not obvious that one can sufficiently easily make a formal distinction between what some humans mean and others mean, we will leave the relation as symmetric—and see how far we get!

It is crucial that the trojan programs exist 'as such' in the particular computer as programs that the user can run: they must have names in an environment. We therefore introduce the concept of a trojan method that characterizes the appropriate constraints. For each type of trojan there will be a different trojan method; but by using the abstraction of a method, we do not consider different representations (i.e. different implementations) of trojans as essentially different. Each trojan method specifies a particular program and a computer configuration supporting an environment in which it can be trojaned. (This pedantry—which is implicit because Ω is 'hidden' in the use of R, E and $[\![\cdot]\!]$—is not only a useful clarification, but reminds one that a trojan of a Unix program sh, say, will not necessarily be a trojan for a different user with a different name space.)

DEFINITION. *A trojan method is a non-empty recursively enumerable relation $T \subseteq R \times R \times L$, such that if $\langle r, \widehat{r}, l \rangle \in T$ then:*[6]

$$\wedge\; r \sim \widehat{r}$$
$$\wedge\; E(r)l \approx E(\widehat{r})l$$
$$\wedge\; \mathcal{M}t \in L^*:$$
$$\wedge\; [\![E(r)l]\!]r \xrightarrow{\;t\;} r'$$
$$\wedge\; [\![E(\widehat{r})l]\!]\widehat{r} \xrightarrow{\;t\;} \widehat{r}'$$
$$\wedge\; r' \not\approx \widehat{r}'.$$

The idea of this is that if $\langle r, \widehat{r}, l \rangle \in T$ for some trojan method T, then $\widehat{r}$ has an environment which is similar to r, but in which the program named l, although looking the same in the two environments if it is executed for most potential inputs, eventually a difference can emerge.

The second line of this definition (i.e. that a trojan does not immediately reveal itself) is optional. The formalism helps make explicit the choices available in the definition of the terms. We feel, however, that it is appropriate, for it is saying a trojan is *initially* indistinguishable from another program but eventually obviously different.

A central contribution of our definition is the notion of a trojan as a relation; however, details in the definition could easily be debated. For example, we could replace the uncertainty of $\mathcal{M}$ by requiring that $\forall s \in L^*$ (i.e. using a *for all* quantifier, rather than the *for most* quantifier) there is an extension t of s with similar properties; the uncertainty has then been pushed into the uncertainty of the length of t. Since trojans generally intend to appear at some point in the future, the majority if not all of them would satisfy this variant definition.

---

[6]We use Lamport's method of writing long formulae [36].

Detection of trojans is built into this definition. A trojan is defined in terms of not being distinguishable from the original (using $\sim$). If a trojan was detectable because it was different it would not be a trojan—it would just be a 'wrong program'.

3.2. Viruses

There are, of course, good viruses and other variations, but we define a virus to be a trojan that additionally infects other named programs, infection being the modification or creation of some program to be a virus. In our framework, then, we do not distinguish a program that *is* a virus and a program that *contains* a virus: to do so would presume an identification of the virus 'code'. (Of course most virus writers write such simple viruses that the distinction has practical use even if no general semantic basis.)[7]

Two representations $r, \widehat{r}$ are virally related on name l if they are part of a trojan method and if the capacity to trojanize and infect is transmitted to other programs. Thus a viral method is a trojan method with added conditions requiring that the method is infectious.

DEFINITION. *A viral method is a trojan method $V \subseteq R \times R \times L$ satisfying the additional condition, such that if $\langle r, \widehat{r}, l \rangle \in V$ then:*

$$\wedge\; \mathcal{M}r_1, r_2 \in R : r_1 \sim r_2$$
$$\wedge\; \exists l' \in (\mathsf{names}\, r_1 \cap \mathsf{names}\, r_2) :$$
$$\langle [\![E(r)l]\!]r_1, [\![E(\widehat{r})l]\!]r_2, l' \rangle \in V.$$

This additional clause is a closure property, saying that evolving two similar representations by virally related programs results in virally related representations. Given a viral method V and a 'normal' representation r, then $\widehat{r}$ is infected by V at l if $\langle r, \widehat{r}, l \rangle \in V$.

It is useful to distinguish an infected system from an infected program, since the cost of establishing whether a system is infection-free is much higher than to establish whether a program is infected.

The definitions do not require a virus to infect with a copy of itself and in particular they allow a virus to encrypt itself in different ways when it infects. Thus we do not require infection to be transitive, since the definition of a virus does not require it to infect with itself (a typical encrypting virus would choose to infect with a differently encrypted variant of itself).

There is nothing in the above definition which requires some syntactic relation to hold between the 'normal' and 'infected' program. This is appropriate, since one could easily imagine a virus incorporating a simple, semantics-preserving re-write system that could be used to transform the combination of the viral code and the new host into some equivalent but syntactically quite different form.

---

[7]Some authors would distinguish a *virus* that only modifies existing programs from a *worm* that can also create programs, typically on another node in a distributed system.

3.3. Summary

An important difference between virus and trojan is now clear: a virus requires the modification of the name space of the representation, thus suitable precautions on naming could inhibit viral spread (under suitable assumptions), whereas a trojan in some sense makes the user do its work and therefore cannot be identified or inhibited if the user is anyway permitted to perform such operations. Trojan compilers form an interesting case where a user may be tricked into performing an evaluation step that can then behave as a virus (Section 4.1.1).

The definitions clearly do not require Ω to be Turing complete in order to support trojan or viral methods. It would be possible for a machine to support programs of viral methods only. Such an apparently limited machine might be of interest for investigations in artificial life [37].

In considering any formal framework of artificial systems, there is always a balance between formalizing what *is* and formalizing what *should be*. Our framework does not handle certain concrete aspects of viruses explicitly: is this a shortcoming in our framework or is it an indication that the complex physical systems that support them should not exist? We think, while it would be a significant achievement to handle more virus behaviour within an elegant framework, it would be a greater achievement to eliminate the possibility of certain sorts of behaviour by better system design.

4. APPLICATIONS OF THE FRAMEWORK

In this framework, where the notion of trojan and virus is inextricably bound up with the environment and the realistic time complexities of testing, the questions one can ask about these phenomena differ from the usual questions. We might consider the following.

- Given a representation r and a viral method V, it is semi-decidable to check whether some other representation $\hat{r}$ is virally related to r in V.
- Given some finite number of infected/non-infected pairs of environments in some unknown viral method V it is impossible to 'infer' V from the data.
- The question, assuming we have to hand a putative virus, 'is p a virus?' makes no sense. For many reasonable notions of $\sim$, even deciding $p \approx \hat{p}$ is undecidable. For very restricted notions of infection (e.g., syntactic modification) limited decidability results are obtainable.
- Is it possible, by elaborating the model of the computing system, to provide a system which resists, detects or is tolerant to viral spread? The affirmative answer changes our attitude to third-party anti-virus products and suggests a requirement, anti-trust notwithstanding, that anti-virus components be integrated into operating systems.
- Following from the previous point: if a (particular) viral method can be recognized, can the representation including it be disinfected, where we take 'disinfected'

to mean some more conservative operation than deletion of all programs overlapping the virus?

- Many programs are constructed out of modules. Our framework does not address this since any collection of modules is just part of the representation. However, in practical terms, there is a difference in convenience or efficiency if we can reliably check modules individually. Most anti-virus products do just this: they normally only check the parts of the representation that are known to change through interaction with the rest of the world—such as when a floppy disc is inserted. The problem does not arise in our framework, but any framework that did model modules (or access rights) would have to be aware of the problem that trojan methods need not be in the modules where they 'should' be—see Section 4.1.1.
- Because anti-virus products are commercial, there are industry tests and league tables. League tables encourage simplistic comparisons, such as 'percentage of wild viruses recognized'. However, hit rates assume a classification of viruses, typically a syntactic one—which arguably inflates the apparent scale of the problem and the efficacy of the anti-virus programs. How should anti-virus products be compared?

4.1. Detectability of trojans

The problem of detecting trojans is at least as hard as determining whether functions are equal, which is undecidable.

There is, of course, a difference between detecting a trojan and resisting the installation of a trojan: security measures are aimed at the latter problem. However, as regards security assumptions precluding an arbitrary program $\hat{p}$ from, in some sense, being related to a program p, by assumption the program $\hat{p}$ is explicitly constructed to trojanize p.

Trojans are not effectively detectable. In fact most trojan and virus detection programs attempt to detect classes of program: the importance of the following result is that many detectors search program representations for patterns (equivalent in biological terms to antigenic epitopes) and attempt to detect any program in the class with that characteristic pattern.

Assuming the classification process is computable and that detection is undecidable, the decidability of detecting classes of trojan would be a contradiction; if the classification process is not computable, then there are trojans that cannot be classified and hence cannot be detected. This has implications for trojan detector programs that attempt to identify specific trojans by their known structure, such as by using a signature.

In many computer environments it is also possible for trojans to dynamically change their code ('mutate'): this would mean that a recently acquired trojan could have a different form to the standard recognized by the trojan detector. By considering the equivalence classes of the behaviours of trojans, we immediately conclude that trojans are not detectable by inspecting their behaviour: this result

is of consequence for so-called gatekeeper detectors that hope to detect trojans or viruses by their (presumably forestalled) actions. They cannot work in general. In practice a useful but insufficient measure of protection is achieved by interpreting primitive operations (such as operating system calls) and intercepting certain computable classes of operation (such as low-level formats); there may be options to enable an operation to proceed if the user deems it to be safe. Inevitably, such methods presuppose a human is capable of making decisions that we have proven undecidable. Inevitably, human mistakes will be made.

Recall that Cassandra, the Trojan prophetess, though correctly warning what the Trojan horse was, was doomed not to be believed!

4.1.1. Thompson's trojan
The intention of Thompson's construction [38] is to use the trapdoor effect of compiling to conceal a trojan from effective detection: $r/c \leftrightarrow r$ is not bijective and r cannot be recovered from r/c (unless c is specially constructed to make this possible). In fact, it may be much worse, there may be no s derivable from $\{c, r/c\}$ such that $r/c = s/c$. This is the well known 'disappearance' of semantics in meta-interpreters (virtual machines) [39]; in Thompson's trojan the semantics that disappear are trojan methods.

Normal compiler bootstrapping is expressed as $c_s/c = c$, where the subscript s conveniently denotes the appropriate source code. Bootstrapping is typically achieved by constructing, by hand or by some other means, an executable program p such that $c_s/p = c_s/\ldots c_s/p$ (it is not necessary that $c_s/p = p$); once p has been successfully applied to c_s, p can be discarded—-although this is risky, as is made clear below. The source code c_s, too, may be discarded or made unavailable (perhaps for commercial reasons). Yet it is still possible to compile all programs. The language c compiles will be complete in some sense (the properties described are not sufficient for Turing completeness).

To follow Thompson's argument it is necessary to introduce conditionals, notated $x \Rightarrow y : z$; we assume that the language processed by c can implement the intended semantics, $x \Rightarrow y : z \stackrel{\text{def}}{=} \textbf{if } x \textbf{ then } y \textbf{ else } z$. It will be sufficient to consider only top-level conditionals and (computable) tests based on identity.

Thompson's discussion is based in C, C compilers, Unix and Unix's login program. We will assume: a non-trivial security-critical program u (say, a login program) and its compiler c, also capable of compiling itself (c_s). We wish to construct a trojan $\widehat{u}$ that is undetectable, *even* given the assumption of the presence of source code u_s of u, which would have provided oracles.

The intention is to leave c_s and u_s unchanged but to have replaced c and u by $\widehat{c}$ and $\widehat{u}$ such that $c_s/\widehat{c} = \widehat{c}$ and $u_s/\widehat{c} = \widehat{u}$, and for $\widehat{c}$ otherwise to behave as c. Once this has been achieved, the trojans will be self-preserving: the trojan method cannot be eliminated easily since everything apart from the behaviour of $\widehat{u}$ will be indistinguishable from normal and it will probably be indistinguishable for 'long enough' from its expected behaviour.

First note that a trojan $\widehat{u}$ of u is easy to detect given u_s, since $\widehat{u} \neq u_s/c$ and we know what u is by applying u_s/c. In other words, with the source u_s we can determine that $\widehat{u}$ is a trojan. In practice one must check all (or only suspect) $u \in \textsf{names}\, r$; however, $\textsf{names}\, r$ is finite and each test is linear.

Suppose now that a compiler c' is constructed, where $s/c' \stackrel{\text{def}}{=} s = u_s \Rightarrow \widehat{u} : s/c$. When applied to u_s, c' trojanizes it to $\widehat{u}$. Note that the test $s = u_s$ is an unproblematic test of identity of representations. Since in all other respects $c' = c$, c can be surreptitiously replaced.

At this stage, there is an undetectable trojan, but the tampering with the compiler is still readily detected since $c_s \neq c'_s$ and $c' \neq c$. The final stage of Thompson's argument removes this loophole.

A further compiler c'' is constructed, where $s/c'' \stackrel{\text{def}}{=} s = u \Rightarrow \widehat{u} : (s = c_s \Rightarrow c'' : s/c)$. This compiler has the remarkable property that it compiles the original c_s to itself, c'', and compiles the target program u_s to a trojan $\widehat{u}$. Since c_s and u_s are original, the new trojan is undetectable. The compiler c'' is bootstrapped as follows.

(1) c''_s is constructed. This is easy, given c_s and the definition of c'' (above).
(2) c''_s is compiled using the original compiler: $c''_s/c \mapsto c''$.
(3) The original compiler's object code c is discarded and replaced by $\widehat{c} = c''$.
(4) The source program c'_s is discarded.

Then the following ensues:

$$s/\widehat{c} = \begin{cases} \widehat{c}, & s = c_s \\ \widehat{u}, & s = u_s \\ s/c, & \text{otherwise.} \end{cases}$$

The source program u_s can now be compiled by $u_s/\widehat{c}$ giving $\widehat{u}$ as required.

We now have a trojan $\widehat{c}$ and a trojan $\widehat{u}$ and no source code other than what is expected, u_s and c_s which have been restored to their originals. All programs compile correctly, except c_s and u_s themselves, but these two cases compile consistently, since the compiler has been trojanized to $\widehat{c} = c''$. The only way to show that $\widehat{u}$ is not in fact u is to find some $r: [\![\widehat{u}]\!]r \neq [\![u]\!]r$—but there is no u available to do this, and even if there was, finding r would be exponentially hard. Login programs such as we have supposed u to be often have features in them specifically to make such operations difficult, since trying to find representations with particular properties has security implications. This trojan will be as difficult to detect as desired by its author.[8]

One can construct a system resistant to the Thompson attack by requiring $s/c \leftrightarrow s$ to be bijective; indeed,

---

[8]Anyone who has bootstrapped compilers will know that discarding independent compilers (the initial p and the subsequent version history of compilers) is foolish: once bugs are introduced—not just deliberate trojans—they can remain even though subsequent compilations remove all signs of them in the source code.

this is readily achieved in a system where representations are directly interpreted and/or there is no direct access to *compiled* forms. Alternatively, denial of access to the source of the compiler is sufficient, with the (awkward) proviso that the source is still required to determine whether the compiler is ever trojaned.

To show that Thompson's construction is sufficient to implement a trojan method, we need to consider his criteria for distinguishability. Having defined this, we must prove that the subverted representation is indistinguishable from the original representation. Thompson's trojan can then be a singleton trojan method provided it guarantees to manifest itself. As Thompson has not specified that it will manifest itself, we come to something of an impasse solved only in that trojans that do not manifest themselves are not a problem!

Thompson distinguishes programs by saying that two programs are indistinguishable if their source code is identical. This sounds constructive, but it is not quite in our framework. So we define two representations to be identical provided the source code of the program under the name of u is identical for both representations. This is poly log computable as the source for u is constant for all representations (given sensible definitions of representations). Unfortunately, this idea fails immediately—the trojan attacks the compiler and you can have identical source code on both representations but different object code. So two representations can be indistinguishable but have different object code for u and therefore be a trojan.

So we add another criterion: the source code for the code for u is the same on both representations and the source code for the compiler is the same for both representations. (This is still poly log.) However, this scheme fails with a trojan that attacks the compiler in a more sophisticated way.

In our terminology, Thompson gave two definitions of $\sim$ and both admitted trojans. The conclusion is not that Thompson trojans are not detectable but that reasonable-sounding definitions of distinguishability do not prevent there being trojans.

The theorem to prove is, given any poly log distinguishing relation, there is a trojan method for that relation. (There has to be some condition like poly log because comparing memory representations without restriction will clearly always show up trojans.) We will take up this challenge in a subsequent paper.

4.2. Detectability of viruses

Is virus detection decidable? If we were to define a relation p virus $\widehat{p}$ just when $\widehat{p}$ is virally related to p for some name l in some viral relation V, we cannot decide the relation because it is at least as hard as function equivalence.

Is virus activity detection decidable? This depends on the computational model assumed, but we can show the infection process can be inhibited under reasonable assumptions. If the environment is fixed, detection is trivial, there being no viruses to detect.

4.2.1. Cohen's proof of virus undetectability

The Cohen proof [23] of the non-computability of detection of viruses is a direct variant of the Halting problem of Turing machines and is therefore subject to the limitations of Turing computability frameworks, as outlined above. The Cohen proof relates to the detection of viruses (i.e. assumed as defined objects), not their methods or behaviour, and it implicitly assumes a fixed Ω. Here we show that this standard proof (widely repeated, for instance in [21, 40] and elsewhere) about the detectability of virus infection is inadequate for a more interesting reason. (In criticizing a proof the intention is to determine exactly what was proved and whether advancements may be made by tightening the proof itself, its assumptions or the theory in which it is expressed.) We quote Cohen's central argument, then discuss a shortcoming:

> In order to determine that a given program P is a virus, it must be determined that P infects other programs. This is undecidable since P could invoke any proposed decision procedure D and infect other programs if and only if D determines that P is not a virus. We conclude that a program that precisely discerns a virus from any other program by examining its appearance is infeasible. In the following [program CV, shown below], we use the hypothetical decision procedure D which returns 'true' iff its argument is a virus, to exemplify the undecidability of virus detection.

```
contradictory-virus :=
{  ...
  main-program :=
  { if ¬ D(contradictory-virus) then
    { infect-executable;
      if trigger-pulled then
        do-damage;
    }
    go next;
  }
}
```

> [...] we have assured that, if the decision procedure D determines CV to be a virus, CV will not infect other programs and thus will not act like a virus. If D determines that CV is not a virus, CV will infect other programs and thus be a virus. Therefore, the hypothetical decision procedure D is self-contradictory, and precise determination of a virus by its appearance is undecidable.

We infer that D does not necessarily evaluate its argument when it attempts to determine whether it is a virus: clearly, to do so would run the risk of activating the virus itself. Cohen implicitly assumes this, since a conventional eager evaluation of his code would abort:[9]

---

[9]In a programming language like Pascal, the parameters of a function call are evaluated before the function can be called. In Cohen's example, this would normally require invoking contradictory-virus.

evaluating `contradictory-virus` would not terminate and indeed would never progress beyond the conditional expression! Instead, D must examine its argument in some safe way which is not specified—the proof is assuming a syntactical representation of a virus. Cohen would like to prove that D cannot work however it is specified.

However, the codes `infect-executable` or `do-damage` are not actually used in the proof and therefore have only rhetorical value in making the program fragment look like a virus. Since, without affecting the proof scheme, any program code (say, `x:=0`) can be substituted (with the corresponding trivial changes to the specification of D) the putative proof is seen to be about the undecidability of program equality—not, more specifically, about virus detection.

We have here, then, an informal proof of a standard result, plus the unjustified assumption that viruses are *modelled* in the formalism of that proof, whereas here they have only been *named*. We agree that to prove that there is no decision procedure one only needs to exhibit a counter example, but we do not agree that `contradictory-virus` is in fact an actual example of a virus. What has happened is that the names `infect-executable` and `do-damage` appeal implicitly to a virus method that may—or may not—be bound to these names *in the computer's representation*. The viral and trojan methods V, T such that

$$\langle\cdot,\cdot,\texttt{infect-executable}\rangle \subseteq V$$
$$\langle\cdot,\cdot,\texttt{do-damage}\rangle \subseteq T$$

are not specified.

4.2.2.   Restricted environments

Viruses require to be able to re-bind identifiers in the environment in order to propagate and cause damage. The obvious solution to this problem is to construct a system which never re-binds names. Though this might seem like a radical proposal it is common in many practical systems.

It is worth making a small technical point here. In most operating systems, a file is bound not to its data (which is what we are modelling), but to where data may be located. In this case, a binding need not change even though the data is changed—for example, text editing a file still leaves it the same file, but with different content. We are deliberately not modelling where data is held and therefore restricting an environment (in our model) to be non-rebinding is an effective restriction on changing the contents of a particularly-named file.

Non-rebinding would require that if $r \xrightarrow{l} r'$, where obviously $l \in \textsf{names}\, r$, then $E(r) \subseteq E(r')$. From this restriction it is immediate that viruses can only infect 'new' name bindings introduced by their execution.

Many task-specific systems such as calculators, dedicated word processors and personal organizers have fixed environments. Even for general purpose computers, many users might be happy to lock the environment so that no names can be rebound; this solution is implemented in a rather crude way in a number of proprietary hardware devices that write protect all of, or parts of discs, though the idea can be generalized [41]. As our framework indicates, though disc locking stops subsequent virus infection, it does nothing to help detect existing infections.

On many computers, the address space of a computer acts as an environment: it maps numbers (very easily generated names!) into representations. Hardware memory protection schemes are practical ways of restricting the address environment so that programs cannot generate names that are mapped into representations in other address spaces. This is quite conventional, but it is a useful example of a restricted environment whose use does not restrict higher-level operations—indeed, the reliability and additional confidence about the behaviour of programs that it confers are highly desirable.

If a system includes a version number in the file names, then no name is ever re-bound, therefore it is impossible for a virus to 'infect' a system unobserved. The user should be able to tidy up the environment from time to time, but this could be a restricted facility requiring direct interaction with the user. The standard advice to users to make backups is no more than a manual (hence unreliable!) procedure to maintain such a non-rebinding environment.

Such a naming proposal seems to go a long way to protecting against viral damage provided the file system functions[10] are effectively virus-proofed. However, this neglects a major component of our framework, namely observation (formalized by the $\sim$ notion that captures the idea that two representations cannot be distinguished by an observer). In many file systems using version numbers, the human interface to the unique names in the file system is constructed to make differences in version number unobservable, for example typically the most recent version of a file will be used by default. In order for the naming scheme to be effective the reverse must be the case—the user must see name changes when they are potentially damaging. This clearly requires major changes in the way user interfaces are constructed.

Turing complete operations on environments (e.g. being able to compute names in an infinite domain) ensure that detection of infection is again undecidable. However, remaining within computability, we can arrange the environment so that certain computations are infeasible without passwords: for example, by using trapdoor functions. The relevance of trapdoors is that (under the appropriate security conditions) the observer and the virus stand on opposite sides of the trapdoor.

Suggested by the framework is the creation and use of names within a system: one can restrict the free use of names usually allowed in current systems. The names in the domain of the environment mapping can be encrypted, but accessed at the observer level via their unencrypted form, thereby making it arbitrarily harder for the virus to find bindings which could be changed unobserved. For example, a programmer writing a program to access a file server demands a key from the environment by supplying a name

---

[10]More precisely, ... functions on R.

and password. This key accesses the binding of that name. Such a scheme would permit programs to compute names (in the encrypted domain of keys), but the probability of computing an unauthorized, but valid, name in the domain of the environment can be limited.

A variety of possible naming schemes might help: indeed it is possible to have an unbounded number of schemes, dynamically created. Various kinds of name servers which impose a management discipline on the creation and use of names could contain the spread of viruses to within any particular set of naming schemes. An obvious application of this is to ensure security when a system changes owner (e.g. when it is first or subsequently sold). A special case is when the naming schemes each contain precisely one binding.

It is often suggested that write-protected executables are immune from infection [42] (who claim that they are immune but impractical). This forgets the environment. If an executable program is unmodifiable that does not imply its access to the environment is unmodifiable: for example, a fixed program may compute names as arguments to the environment. A virus could therefore change the behaviour of a program by affecting this computation (e.g. by initializing it with different data). A realistic case of this situation is that of a program that runs some background server or other process: it computes a name to access the environment (e.g., in the simplest case by reading a data file of server names) to load the background program, but a virus might simply cause it to load the wrong program.

The what-might-be-called 'the write-protected executable fallacy', that one is supposedly safe when executable programs are write protected, confuses the security of the program for the security of the environment.

4.2.3. Viral resistance

In the practical use of a computer, the user only observes some of the outputs of a computation and only provides some of its inputs. The problem of viruses is that they are primarily concerned with inputs and outputs that the user normally ignores at the time of the computation. For example, the program code itself is not normally considered one of its own inputs, but this is precisely where most viruses reside and how they directly affect the input of the program; a virus's output may alter files containing other programs, of which the user is unaware.

A virally resistant system can be constructed by introducing observations O, which are to be communicated to the user. We extend $E: R \rightarrow (L \twoheadrightarrow R \times O)$ and $[\![\cdot]\!]: R \rightarrow (R \rightarrow R) \times (R \rightarrow O)$. Names are now bound to pairs $\langle p, o \rangle$ and the meaning of the pair is a pair of functions, one of which computes the result of doing the command and the other 'observes' the result to see it passes some checks. Observes, in the sense we are using it, means 'prepared in some way that can be tested by a user'.

In running a program $\langle p, o \rangle$ the system runs p as before to obtain the results and the new environment and runs o to observe the result, presenting the observation to the user. Programs have lots of inputs and results over the representation space, but a user's tests do not explore the domain uniformly, being only interested in conventional inputs—likewise, they only examine conventional outputs, say on the screen, not outputs that change files. The component o makes it clear that all the output must be observed.

By definition, a virus changes a binding of some name from $\langle p, o \rangle$ in the environment to some new $\langle \widehat{p}, \widehat{o} \rangle$. In general, it is clearly not possible to compute $\widehat{o}$ from $\langle p, o \rangle$ to ensure that in an arbitrary environment $\widehat{o}$ computes the same value after a run of $\widehat{p}$ as o does after a run of p. The value of o must be interpreted by the observer; it is insufficient for o to yield a specific token (say true) for any authenticated binding, since any predetermined token can easily be computed by $\widehat{o}$. Thus given some notion of an external observer (e.g. the user) eventually any virus can be detected. Astute choices of o and observer make the likelihood of prompt detection much higher—the observer can be hardware (the range of o can be digital signatures).

A more intriguing idea is for the result of o to be a pattern (e.g. a video bitmap) and to rely on the human skill of recognizing patterns and changes in patterns [43]—maybe a virus would show up as, say, an irritating line across the screen. This is an attempt at distribution free testing [44], aided by human sense perception. Distribution free testing is a more mechanical process that systematically samples the input/outputs so that o gives a 'fair' (distribution free) sample of the program's complete effect, though doing this depends on some properties of the program, but does *not* depend on knowing what the correct output should be. (Good cryptographic detection techniques are to some extent attempts to find suitable distribution free sampling functions.) Finally, so that it cannot be compromised, o may be implemented by hardware.

Implementations of such schemes must be undertaken very carefully and some obvious implementations are suspect, simply because an implementation that (say) provides programs as pairs $(E \rightarrow R \times E) \times (E \rightarrow O)$ may accidentally provide operations that compromise the system. Thus, an unadorned Turing machine can readily implement this scheme, but does not ensure that access functions for the pairs $\langle \widehat{p}, \widehat{o} \rangle$ are excluded: which, of course, would defeat the whole object of the distinction—it would be possible to construct an $\widehat{o}$ that simply replayed the output of o. See Section 4.3 for further discussion.

The invention of asymmetric (public key) cryptography challenged a central assumption, that security could be achieved through secrecy and obscurity. (Equally, the secrecy could conceal incompetence.) Now, new cryptographic algorithms have been widely published and widely scrutinized [45], and this scrutiny increases confidence in their effectiveness. It is interesting, then, to note that many discussions of viruses (e.g. [46]) do not wish to reveal anti-virus methods. Perhaps we need a similar breakthrough in viral resistance?

4.3. Virtual machines

Many programs (such as spreadsheets, language interpreters like TEX and commercial word processors)[11] introduce virtual machine environments. These virtual machines may be 'vectors' for infecting with viruses even though they run on otherwise secure operating systems. Virtual machine environments overcome attempts at protecting the implementation machine.

Since some programs written in a system $\mathcal{L}$ (BASIC, Java, Microsoft Word macros, . . .) need to (say) delete files or have other permissions, then $\mathcal{L}$ needs those capabilities. An $\mathcal{L}$ system runs on a system which may itself be protected from virus activity, but the $\mathcal{L}$ system creates an environment for running $\mathcal{L}$ programs. This not only enables rebindings but changes the observed behaviour of the computer—of course, it must, since one wants to run the system $\mathcal{L}$! Thus $\mathcal{L}$ creates a virtual machine: an $\mathcal{L}$-machine simulated by the PC-machine. Clearly, our framework applies at each level of virtual machine and this has significant repercussions for the sort of virtual machine one would like to support in a secure system. In particular, the user interface must make an *observable* distinction between each virtual machine (otherwise they could alias each other). Even in Java, which is designed with networked programming in mind, this distinction is made by libraries, not intrinsically.

The severity of the problem introduced by virtual machines is shown by Thompson's demonstration that explicit code (including code containing viral methods) in a virtual machine can be made to disappear from one level of the virtual machine by embedding it in an appropriate form in the implementation machine (Section 4.1.1). If the virtual machine supported is Turing complete and supports actions such as modifying the environment map (e.g. by permitting writing to files), then it is not possible to detect viral methods. All 'useful' virtual machines meet these two conditions.

4.4. A note on object-orientation

The increasing popularity of object-oriented programming and icon-based user interfaces (where there are very many observable objects in the environment) is based on claims on their efficiency and convenience of programming. Although the run-time systems of object-oriented systems (Java being an example) may take steps to be secure, object-orientation itself is at odds with secure computation. To the extent that object-orientation has an impact on programmer convenience, it is clearly dependent on large numbers of computationally simple bindings. Inheritance is a problem because it provides a recursive environment. Indeed, Java has recently suffered from the `StrangeBrew` virus, which infects the Java environment—and Java, being platform independent, ensures that the virus can run on almost any type of computer [47].

In systems that have inheritance, operations have default behaviour. Bontchev [46] gives several concrete examples based on a macro language. We give an abstract characterization of one of his examples: suppose there is an easily recognized virus consisting of a set of macros, S. (Typically, one of the components will be activated by user activity, such as opening a file, and on running it will install another component as its payload.) A virus writer modifies S to make a variant. Now anti-virus software may recognize only the original components of this new virus and eliminate them; however what remains may be an intact virus because the 'missing' components inherit default implementations. Ironically, this third, new, virus was created by the anti-virus procedure!

A thorough analysis of these issues is beyond the scope of this paper, except to note that any correct formal computation expressible in an object-oriented paradigm is expressible in another, possibly more secure, paradigm—but the real issue here is actually the trade-off between observable properties, the relationships of names in the environment and other aspects of usability and security.

5. KOCH'S POSTULATES

Robert Koch, the distinguished bacteriologist, contributed four criteria, known as Koch's Postulates, for identifying the causative agent of a particular biological disease.

(1) The pathogen must be present in all cases of the disease.
(2) The pathogen can be isolated from the host and grown in pure culture.
(3) The pathogen from the culture must cause the disease when inoculated into a healthy, susceptible host.
(4) The pathogen must be isolated from the new host and shown to be the same as the original.

To make sense of Koch's Postulates in our framework we may equate *pathogen* with *viral method*. It follows that a biological-type 'isolation' (Postulate 2) is non-computable. To the extent, then, that Koch's Postulates capture biological criteria, biological metaphors cannot be applied with any fidelity to computer virus phenomena. Possibly Koch would have had a different view if biological pathogens were better able to mutate rapidly and maintain their (viral) method.[12] Because biological pathogens do not do this, Koch's Postulates can be usefully expressed with respect to representations rather than interpretations. A more appropriate biological metaphor for computer viruses is Dawkins's *meme* [26], for this corresponds to a software configuration running in the virtual machine provided by the hardware of a brain. (Dawkins makes explicit the connection with computer viruses.)

Given current interest in prions and transgenic infections (e.g. via xenotransplants) a formal framework for biological applications would be desirable. The way in which semantics in meta-interpreters disappears (exploited in Thompson's trojan) obviously has profound implications

---

[11]Some authors call such viruses *macroviruses*; however, the viral methods typical of macroviruses (see [46]) are not restricted to macro languages *per se*. We suggest this terminology is misleading.

[12]Biological viruses mutate rapidly (in biological terms) but do not evolve rapidly [48].

and may help understand prions. In any case, such results would certainly apply to replication using DNA. Unfortunately, our framework makes certain assumptions that apply specifically to what might be called typical electronic digital computers: whilst certain sorts of computation can be performed to order (e.g. to detect infection), one is not prepared to devote excessive resources to this. In a biological context, the resources available and how they can be recruited are very different. Immune systems are massively parallel and autonomous, yet they are very slow to produce new antigens (vaccination is a rehearsal for the immune system). Biological replication, whilst comparatively slow, occurs in parallel at a molecular or species level but serially at an individual level. Computer viruses typically do not mutate using genetic algorithms, but rather use more specialized techniques (e.g. encryption) that guarantee viable replication. Thus there are significant differences, which are beyond the scope of this paper to explore satisfactorily.

Notwithstanding the fundamental biological differences, there is of course promise in biologically-inspired techniques for detecting and fighting viruses. See [9] for an insightful general discussion and [10] which describes a prototype 'digital immune system'. (Coincidentally, the preceding article in the same journal gives examples of biological viruses that successfully suppress their hosts' immune systems [49]!)

6. CONCLUSIONS

A new framework has been introduced that appears to be better than previous attempts at addressing trojan and viral issues. Its main merit is that it is productive in raising and helping clarify the sorts of issues that need addressing. Although it abstracts away from the richness of the phenomena, it accounts for most of the concrete features: it makes clear that viruses are a very complex notion—involving the naming of objects, their behaviour and the observation of that behaviour.

Our framework for computer virus infection shows that Koch's Postulates are inadequate for the phenomenon of computer viruses; in other words, the medical/biological metaphor for computer virus behaviour is seriously misleading.

A virus is a program that, in addition to having a trojan activity, infects other programs. We have shown that a Turing machine equivalent model is insufficient to capture important details of virus behaviour. As contributions towards a theory of computer viruses we pointed out that formalism as such has no notion of expected behaviour, against which undesirable behaviour can be compared. Infection is with respect to an environment and must be identified by an observer using finitary tests. It follows that suitable constraints on environment operations can inhibit both trojan and virus infection.

We have given a proof that trojan code in general cannot be detected. Classes of trojan cannot be detected either and this result puts limits on what can be expected of both pattern-matching type detectors and detectors that rely on intercepting certain sorts of behaviour. We have suggested various forms of observation as appropriate to control viruses.

We have shown that virus infection can be detected and limited. It follows that the spreading of viral methods can be restricted, but once infected by a virus there are limitations on what can be done to detect it, either by its unwanted behaviour, its code signature or any other characteristic. Whether users of computers would wish to convert to a new architecture more secure against infection is a question we do not address here; necessarily such computers would be incompatible with existing computers [50]—merely being discless network computers will not be sufficient.

Finally, we admit we are not yet satisfied. Although we have introduced and motivated important distinctions, the framework itself is unwieldy and the distinctions are hard to maintain in applied reasoning. It is hard to derive interesting theorems. Nevertheless we have successfully shown that viruses are a very complex phenomenon, despite frequently exhibiting such utterly banal behaviour that we would rather dispel them from our minds—if not just from our computers. Just as the current variety of viruses is not the last word in deviousness, our framework is not the last word in theoretical work with computer viruses. We hope our lasting contribution will be a greater awareness amongst system designers of the possibilities that unnecessarily liberal programming environments provide hackers. We hope, too, to have challenged other theorists to pursue some of the interesting and important formal questions begged by taking viruses seriously.

REFERENCES

[1] Brunner, J. (1993) Sometime in the recent future .... *New Scientist*, **138**, 28–31.

[2] Wegner, P. (1997) Why interaction is more powerful than algorithms. *Commun. ACM*, **40**, 80–91.

[3] Meinel, C. P. (1998) How hackers break in . . . and how they are caught. *Sci. Amer.*, **279**, 70–77.

[4] Bissett, A. and Shipton, G. (1998) Envy and destructiveness: understanding the impulses behind computer viruses. In *Proc. 4th Int. Conf. on Ethical Issues in Information Technology, Ethicomp'98*, pp. 97–108.

[5] Virgil (19BC) *The Aeneid*, Book II.

[6] Anderson, J. P. (1972) *Computer Security Technology Planning Study*, ESD-TR-73-51, vols **I** & **II**. USAF Electronics Systems Division, Bedford, MA.

[7] Cohen, F. (1994) *A Short Course On Computer Viruses* (2nd edn). John Wiley, New York.

[8] Denning, D. E. R. (1983) *Cryptography And Data Security*. Addison-Wesley, Reading, MA.

[9] Forrest, S., Hofmeyr, S. A. and Somayaji, A. (1997) Computer immunology. *Commun. ACM*, **40**, 88–96.

[10] Kephart, J. O., Sorkin, G. B., Chess, D. M. and White, S. R. (1997) Fighting computer viruses. *Sci. Amer.*, **277**, 88–93.

[11] Spafford, E. H. (1994) Computer viruses as artificial life. *Artificial Life*, **1**, 249–265.

[12] *Virus Bulletin*, ISSN 0956–9979.
URL: `http://www.virusbtn.com/`

[13] Hoffman, L. J. (1990) *Rogue Programs: Viruses, Worms And Trojan Horses*, p xi. Van Nostrand Reinhold, New York.

[14] Stevens, M. (1998) Pest control. *New Scientist*, **159**, 64.

[15] Goodenough, O. R. and Dawkins, R. (1994) The 'St Jude' mind virus. *Nature*, **371**, 23–24.

[16] Rhodes, R. (1994) Chain mail. *Nature*, **372**, 230.

[17] Jones, S. K. and White, C. E. Jr. (1990) The IPM model of computer virus management. *Comput. Security*, **9**, 411–418.

[18] Turing, A. M. (1939) Systems of logic based on ordinals. *Proc. London Math. Soc.*, Series 2, **45**, 161–228.

[19] Witten, I. H., Thimbleby, H. W., Coulouris, G. F. and Greenberg, S. (1991) Liveware: a new approach to sharing data in social networks. *Int. J. Man-Machine Studies*, **34**, 337–348.

[20] Coulouris, G. F. and Dollimore, J. (1988) *Distributed Systems*. Addison-Wesley, Reading, MA.

[21] Burger, R. (1989) *Computer Viruses, A High-tech Disease* (3rd edn). Abacus, Data Becker, Düsseldorf.

[22] Stevens, K. (1994) Mind control. *Nature*, **372**, 734.

[23] Cohen, F. (1987) Computer viruses. *Comput. Security*, **6**, 22–35.

[24] Ferbrache, D. (1992) *A Pathology Of Computer Viruses*. Springer, London.

[25] Nachenberg, C. (1997) Computer virus–antivirus coevolution. *Commun. ACM*, **40**, 46–51.

[26] Dawkins, R. (1989) *The Selfish Gene* (2nd edn). Oxford University Press, Oxford, 1989.

[27] Adleman, L. M. (1988) An abstract theory of computer viruses. In Goldwasser, S. (ed.), *Advances in Cryptology—CRYPTO'88, (Lecture Notes in Computer Science,* **403***)*, pp. 354–374. Springer, Berlin.

[28] Cohen, F. (1989) Computational aspects of computer viruses. *Comput. Security*, **8**, 325–344.

[29] Thimbleby, H. W. and Anderson, S. O. (1990) 'Virus theory', Institution of Electrical Engineers Colloquium. In *Viruses and Their Impact on Future Computing Systems*, pp. 4/1–4/5. Institution of Electrical Engineers Publication No. 1990/132.

[30] Thimbleby, H. W., Witten, I. H. and Pullinger, D. J. (1995) Concepts of cooperation in artficial life. *IEEE Trans. Syst., Man Cyber.*, **25**, 1166–1171.

[31] Ladkin, P. B. and Thimbleby, H. W. (1994) Comments on a paper by Voas, Payne and Cohen, 'A model for detecting the existence of software corruption in real time'. *Comput. Security*, **13**, 527–531.

[32] Bates, J. (1990) WHALE . . . a dinosaur heading for extinction. *Virus Bulletin*, November 17–19. See [12].

[33] Thimbleby, H. W. (1987) Optimizing self-replicating programs. *Comput. J.*, **30**, 475–476.

[34] Fokkinga, M. (1996) Expressions that talk about themselves. *Comput. J.*, **39**, 408–412.

[35] Kanada, Y. (1997) Web pages that reproduce themselves by Javascript. *ACM SIGPLAN Notices*, **32**, 49–56.

[36] Lamport, L. (1994) How to write a long formula. *Formal Aspects Comput.*, **6**, 580–584.

[37] Langton, C. (1988) Artificial life. In Langton, C. (ed.), *Artificial Life (Santa Fe Inst. Studies in the Sciences of Complexity)*, pp. 1–47. Addison-Wesley, Reading, MA.

[38] Thompson, K. (1987) Reflections on trusting trust. *ACM Turing Award Lectures*. In Ashenhurst, R. L. and Graham, S. (eds), *ACM Turing Award Lectures*, pp. 171–177. Addison-Wesley, Reading, MA.

[39] Henderson, P. (1980) *Functional Programming*. Prentice-Hall, Englewood Cliffs, NJ.

[40] Leiss, E. L. (1990) *Software Under Siege*. Elsevier Science Publishers, Oxford.

[41] Thimbleby, H. W. (1991) An organizational solution to piracy and viruses. *J. Syst. Software*, **25**, 207–215.

[42] Pozzo, M. M. and Gray, T. E. (1987) An approach to containing computer viruses. *Comput. Security*, **6**, 321–331.

[43] Race, J. (1990) Using computer graphics to find interesting properties in data. *Comput. Bull.*, IV, **2**, 15–16.

[44] Lipton, R. J. (1991) New directions in testing. In Feigenbaum, J. and Merritt, M. (eds), *Proc. DIMACS Workshop in Distributed Computing and Cryptography (DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 2)*, pp. 191–202.

[45] Zimmermann, P. (1995) *PGP: Source Code And Internals*, MIT Press, MA.

[46] Bontchev, V. (1998) Macro virus identification problems. *Comput. Security*, **17**, 69–89.

[47] Hancock, B. (1998) Security views (Java gets a foul taste—first reputed Java virus). *Comput. Security*, **17**, 462–474.

[48] Huang, A. S. and Coffin, J. M. (1992) Virology: how does variation count? *Nature*, **359**, 107–108.

[49] Beckage, N. E. (1997) The parasitic wasp's secret weapon. *Sci. Amer.*, **277**, 82–87.

[50] Thimbleby, H. W. (1991) Can viruses ever be useful? *Comput. Security*, **10**, 111–114.