

INFORMATION TRANSMISSION IN COMPUTATIONAL SYSTEMS

Ellis Cohen
University of Newcastle upon Tyne

This paper presents Strong Dependency, a formalism based on an information theoretic approach to information transmission in computational systems. Using the formalism, we show how the imposition of initial constraints reduces variety in a system, eliminating undesirable information paths. In this way, protection problems, such as the Confinement Problem, may be solved. A variety of inductive techniques are developed useful for proving that such solutions are correct.

Section 1. INTRODUCTION

Users of large computational systems require assurances that information they deem private remains private. Recent work (e.g. [Denning 76, Walters 75]) has used a notion of "information flow" in order to formally state such assurances. If, for each operation provided at some level of a system, one can determine the information flow resulting from that operation - inductive techniques can be used to determine the extent of information transmission in the system as a whole.

The information flow resulting from execution of a single operation can formally be determined from the semantics of the operation. However, except for work by Jones & Lipton [75], operational information flow has been based on the syntax of the operation instead [Denning and Denning 76, Millen 76].

This paper introduces Strong Dependency, which derives the information flow from the functional semantics of an operation. (A companion paper [Cohen 77] derives axiomatic techniques for determining the operational information flow from syntax.) Furthermore, it provides a unified theory of information transmission for both single operations and for the entire system.

Strong Dependency is based on ideas taken from classical information theory [Ashby 56]. Notably, it borrows the idea that information is transmitted from a source to a destination only when variety in the source can be conveyed to the destination. Sections 2-4 provide background definitions and formalise these ideas.

The research described in this paper was supported at Carnegie-Mellon University by the Defense Advanced Research Projects Agency (F44620-73-C-0074) where it is monitored by the Air Force Office of Scientific Research, and by the National Science Foundation under grant MCS75-07251A01.

Initial constraints on a system can decrease the variety in a system, reducing the variety conveyed and preventing transmission of information. Initial constraints are important, for current systems tend to provide fairly elaborate mechanisms [Cohen and Jefferson 75] for preventing information transmission, however, these must be initialised properly in order to have the desired effect. Sections 5 and 6 show how the Strong Dependency formalism can be used to state what must be proven - the absence of certain information transmission paths given proper initialisation.

Proving the absence of information transmission over all possible executions of a system requires an inductive proof technique. Section 7 accordingly derives Strong Dependency Induction, which is shown to be a generalisation of the information flow induction techniques mentioned above. Section 8 shows how Strong Dependency Induction might be used, by certifying a solution to the Confinement Problem [Lampson 73] in a simple system.

Information flow has generally been defined as transitive. That is, if information can flow from a to m and from m to b, then a flow from a to b is assumed as well. Section 9 illustrates two systems where transitivity fairly obviously does not hold. Moreover, a technique for Separation of Variety is derived, which can be used in conjunction with Strong Dependency Induction to prove absence of information transmission in non-transitive cases.

Section 2. COMPUTATIONAL SYSTEMS

A Computational System is formally specified as a triple $\langle NM, \Sigma, \Delta \rangle$. We may imagine that a system is comprised of a set of objects (e.g. variables, files, processes) named by NM. The value (contents) of some object x ($x \in NM$) depends upon the state of the system σ ($\sigma \in \Sigma$). I write $\sigma.x$ to mean the value of x in state σ .

Execution of an operation alters the state. Formally we define an operation δ ($\delta \in \Delta$) as a mapping from states to states and interpret $\delta(\sigma) = \sigma'$ to mean that σ' is the new state after δ is executed in state σ . Informally we can define operations by any convenient programming-language-like notation. For example, if σ were some state such that $\sigma.a = 11$, and δ were the operation $\delta: b \leftarrow a$, then the state $\delta(\sigma)$ would be identical to σ except that b 's value in state $\delta(\sigma)$ would be 11 ($\delta(\sigma).b = 11$), regardless of its value in state σ .

I write $\sigma \stackrel{\bar{x}}{=} \sigma'$ to mean that σ and σ' must have the same values for all objects except x . Formally

$$\sigma \stackrel{\bar{x}}{=} \sigma' \equiv_{\text{def}} (\forall y \neq x) (\sigma.y = \sigma'.y)$$

In the example above, $\delta(\sigma) \stackrel{\bar{b}}{=} \sigma$.

A history is defined as a sequence of operations. If H were $\delta_1\delta_2\delta_3$ then execution of H indicates sequential execution of δ_1 , δ_2 and then δ_3 .

$$H(\sigma) = (\delta_1\delta_2\delta_3)(\sigma) = \delta_3(\delta_2(\delta_1(\sigma)))$$

Section 3. INFORMATION TRANSMISSION

Information transmission in computational systems can be formalised by considering the classical treatment of information theory [Ashby 56]. For example, consider execution of

$$\delta: b \leftarrow a$$

with respect to an information channel having a as a source and b as a destination. Our task is to determine whether information is transmitted over that channel. Classical information theory tells us that

Information is transmitted over a channel when variety is conveyed from the source to the destination

For example, if (in a 16 bit machine) a could (with equal probability) take on any of 2^{16} possible values, then there are $\log_2(2^{16}) = 16$ bits of variety in a and execution of δ conveys all of that variety to b , for after execution of δ , b may also take on 2^{16} values, each one corresponding to an initial value of a .

In general, information can be transmitted from a to b over execution of H if, by suitably varying the initial value of a (exploring the variety in a), the resulting value in b after H 's execution will also vary (showing that the variety is conveyed to b). Note that this formulation implies that if a is known to be a particular constant, then no information can be transmitted from a to b . If a initially contains no variety, then surely none can be conveyed to b .

Section 4. STRONG DEPENDENCY

Classical information theory is concerned with the amount of information transmitted over a channel. Strong Dependency is only concerned with whether any information can be transmitted at all. I write $a \stackrel{H}{\triangleright} b$, b Strongly Depends on a over execution of H , to mean that variety in a can be conveyed to b over execution of the sequence of operations H . Formally

$$a \stackrel{H}{\triangleright} b \equiv_{\text{def}} (\exists \sigma_1 \sigma_2) (\sigma_1 \stackrel{\bar{a}}{=} \sigma_2 \wedge H(\sigma_1).b \neq H(\sigma_2).b)$$

$\sigma_1 \stackrel{\bar{a}}{=} \sigma_2$ represents two suitably chosen states that differ at a alone - i.e. exhibiting variety in a . That variety can be conveyed to b , if after execution of H , the value of b differs in the two states.

Figure 1 presents two examples of Strong Dependency, both showing how information can be transmitted from a to b over execution of a single operation.

Example 1: $\delta: b \leftarrow a$

	σ_1	σ_2	$\delta(\sigma_1)$	$\delta(\sigma_2)$
a	10	20	10	20
b	40	40	10	20

Example 2: $\delta: \text{if } a \text{ then } b \leftarrow x$

	σ_1	σ_2	$\delta(\sigma_1)$	$\delta(\sigma_2)$
a	true	false	true	false
b	40	40	60	40
x	60	60	60	60

$$\sigma_1 \stackrel{\bar{a}}{=} \sigma_2 \quad \delta(\sigma_1).b \neq \delta(\sigma_2).b$$



Figure 1

Example two has intrigued previous researchers, for if information transmission is perceived as a "flow" of information, then when a is false, as in σ_2 , execution of δ has no effect and can hardly result in a "flow" of information from a to b . Yet the fact that b 's value does not change, of course indicates that a must be false, representing new information gained about a . Jones and Lipton [75] and Denning [76] discuss these situations using the terms "negative inference" and "implicit flow" respectively. These cases need not be treated specially. Using Strong Dependency, we never ask whether there is a "flow" of information when a history is executed in some particular state; rather we compare the effect of execution in different states. Analogously, in classical information theory, the information transmitted over a channel is not determined by the contents of any single message, but rather by the variety in the set of contents of all possible messages.

Section 5. TRANSMISSION AND CONSTRAINTS

This section considers the effect of constraining a system so that it is permitted to begin execution only in some subset of all possible initial

states. The approach to constraint again derives from classical information theory:

Constraint reduces the variety in a system potentially reducing the variety that may be conveyed

For example, if a system is initially constrained so that a can only take on values from \emptyset to 7, then execution of

$\delta: b \leftarrow a$

can still transmit information from a to b. This is written

$a \stackrel{\delta}{\Downarrow}_{\phi} b$ where

$\phi(\sigma) \equiv \emptyset \leq \sigma.a \leq 7$

However if δ' is the operation

$\delta': b \leftarrow a \text{ div } 20$ (integer division)

then, the resulting value of b is always \emptyset when ϕ is satisfied initially. No variety is conveyed so imposing ϕ prevents information transmission from a to b and

$\neg a \stackrel{\delta'}{\Downarrow}_{\phi} b$ even though $a \stackrel{\delta'}{\Downarrow}_{\phi} b$

The definition of Strong Dependency determined the initial variety in a by examining arbitrary pairs of states that differed at a alone. In the presence of the constraint ϕ , we need only consider states that satisfy ϕ . Formally

$\sigma_1 \stackrel{\phi}{\equiv} \sigma_2 \equiv_{\text{def}} \phi(\sigma_1) \wedge \sigma_1 \stackrel{x}{=} \sigma_2 \wedge \phi(\sigma_2)$

$a \stackrel{H}{\Downarrow}_{\phi} b \equiv_{\text{def}} (\exists \sigma_1 \sigma_2) (\sigma_1 \stackrel{\phi}{\equiv} \sigma_2 \wedge H(\sigma_1).b \neq H(\sigma_2).b)$

The examples above showed how information transmission could be prevented by reducing the variety in the source a. A more traditional sort of constraint has been discussed previously by Millen [76], who notes that execution of

$\delta: \text{if } q \text{ then } b \leftarrow a$

causes no transmission of information from a to b if q is constrained to be false. Formally, it can be shown that

$\neg a \stackrel{\delta}{\Downarrow}_{\phi} b$ where

$\phi(\sigma) \equiv \neg \sigma.q$

I've implied that $\neg a \stackrel{H}{\Downarrow}_{\phi} b$ guarantees that information cannot be transmitted from a to b over transmission of H given any ϕ . In fact, this guarantee may not hold for constraints that are not autonomous. A formal definition of autonomy and additional details can be found in [Cohen 76]. Informally, ϕ is autonomous if the constraints on the values of each object are independent.

Examples of autonomous constraints are:

$\phi(\sigma) \equiv \sigma.a \neq 7 \wedge \sigma.b = 12$
 $\phi(\sigma) \equiv (\forall x)(\sigma.x \leq 18)$

Non autonomous constraints are:

$\phi(\sigma) \equiv \sigma.a \leq \sigma.b$
 $\phi(\sigma) \equiv (\forall a,b)(b \in \sigma.a.\text{refs} \supset \sigma.a.\text{level} = \sigma.b.\text{level})$

where refs and level are distinct subcomponents of state σ . This last constraint can be read as: If b is in the set of objects a can reference in state σ , then the (dynamic security) level of a and b are the same. The dynamic nature of levels requires a non-autonomous constraint. If security levels were statically associated with objects, level(a) indicating a's level, the constraint would be autonomous.

$\phi(\sigma) \equiv (\forall a,b)(b \in \sigma.a.\text{refs} \supset \text{level}(a) = \text{level}(b))$

The value of a in state σ is dependent upon the properties of other objects (i.e. level(b)), but independent of their value. The remaining examples in this paper will all use autonomous constraints. (See the Appendix for effects of non-autonomy.)

Section 6. INFORMATION PROBLEMS AND SOLUTIONS

This section discusses how the Strong Dependency formalism can be used to state two well known problems - the Security Level Problem [Denning 76] and the Confinement Problem [Lampson 73] in a manner that is independent of any particular computational system.

The last section discussed the operation

$\delta: \text{if } q \text{ then } b \leftarrow a$

and showed that although $a \stackrel{\delta}{\Downarrow}_{\phi} b$, imposition of the constraint

$\phi(\sigma) \equiv \neg \sigma.q$

prevented transmission of information from a to b. We can think of ϕ as solving the information problem

$\neg a \stackrel{\delta}{\Downarrow}_{\phi} b$

That is - guarantee that no information may be transmitted from a to b. A solution to the problem is represented by a constraint ϕ whose imposition prevents the specified information transmission. ϕ can be thought of as characterizing what are often referred to as "secure states".

The Confinement problem is concerned with guaranteeing the privacy of information held by individual users. The problem is typically described in terms of the difficult situation where a user engages a service. The private information must be passed to that service and so seems to be temporarily out of direct control of

the user. A number of solutions to this problem have been proposed [Rotenberg 73, Cohen & Jefferson 75, Lampson 73]. This paper will concentrate on a more general formulation of the problem instead.

Objects (including executors) can be divided into three groups. Those characterised as "Private" corresponding to objects directly under the control of the owners of the information. "Spy" characterises the objects to which that information must not be transmitted. The remainder of the objects can temporarily hold private information as long as it is not subsequently transmitted to a spy.

The confinement problem can then be specified as

$$(\forall x,y) (x \Downarrow_{\phi} y \supset \text{Private}(x) \supset \neg \text{Spy}(y))$$

$$\text{where } x \Downarrow_{\phi} y \equiv_{\text{def}} (\exists H) (x \Downarrow_{\phi}^H y)$$

That is, if execution of any history transmits information from a private object to some other object, then that other object must not be a spy.

Of course, each user of a system may define confinement differently, in terms of different sets of "private" and "spy" objects; each one defining a different instance of the Confinement problem. Each instance of the problem determines a different set of solutions ϕ . Thus a system itself does not solve the Confinement problem. Rather a system may provide mechanisms that permit a user to produce secure states and which then guarantees that subsequent states remain secure.

The Security level problem is usually described as having less of an individualistic flavour. Each object x is initially assigned a security level, $\text{Level}(x)$ determined on a system wide basis. No information is to be transmitted from an object at one level to an object at a lower level. Formally

$$(\forall x,y) (x \Downarrow_{\phi} y \supset \text{Level}(x) \leq \text{Level}(y))$$

Section 7. STRONG DEPENDENCY INDUCTION

Proving the correctness of solutions to the Confinement problem or the Security level problem requires an analysis of information paths over all possible histories. Strong Dependency Induction makes that task practical. Intuitively, it is based on the notion that if information is transmitted from a to b over execution of HH' , then there must be some intermediate object m (perhaps the same as a or b in degenerate cases) such that execution of H transmits information from a to m and execution of H' transmits information from m to b . Formally it can be proven that the following property holds.

PROPERTY 1 (Flow Property)

$$a \Downarrow^{HH'} b \supset (\exists m) (a \Downarrow^H m \wedge m \Downarrow^{H'} b)$$

When an initial constraint ϕ is imposed on the system, the theorem continues to hold as long as ϕ is both invariant (that is $(\forall \sigma, \delta) (\phi(\sigma) \supset \phi(\delta(\sigma)))$) and autonomous (other generalisations may be found in [Cohen 76]).

PROPERTY 2 (Flow with Constraint)

If ϕ is invariant and autonomous then

$$a \Downarrow_{\phi}^{HH'} b \supset (\exists m) (a \Downarrow_{\phi}^H m \wedge m \Downarrow_{\phi}^{H'} b)$$

The following property follows directly.

PROPERTY 3 (End point Isolation)

If ϕ is invariant and autonomous then:

To show $\neg a \Downarrow_{\phi} b$ prove

$$(\forall m \neq b, \delta) (\neg m \Downarrow_{\phi}^{\delta} b) \text{ or}$$

$$(\forall m \neq a, \delta) (\neg a \Downarrow_{\phi}^{\delta} m)$$

That is, if no single operation can transmit information to b from any other object or if no single operation can transmit information from a to any other object, then information cannot be transmitted from a to b over any history.

More generally, we may want to show that $x \Downarrow_{\phi} y$ can be permitted, but only for certain classes of $\langle x,y \rangle$ pairs, for example when $\text{Level}(x) \leq \text{Level}(y)$. That relationship is both reflexive and transitive, and in general

PROPERTY 4 (Lattice Induction)

If ϕ is invariant and autonomous and q is reflexive and transitive then

$$\text{To show } (\forall x,y) (x \Downarrow_{\phi} y \supset q(x,y))$$

$$\text{Prove } (\forall x,y,\delta) (x \Downarrow_{\phi}^{\delta} y \supset q(x,y))$$

Thus for the Security level problem, it is only necessary to show that if x 's security level is greater than y 's, no single operation transmits information from x to y .

Reflexive, transitive relations arise naturally in the definition of information problems. As a result, a number of researchers [Denning 75, Walters 75, Jones & Lipton 76] have suggested an information flow relation that is itself reflexive and transitive. Such a relation is analogous to Strong Dependency having the property that

$$a \Downarrow^H m \wedge m \Downarrow^{H'} b \supset a \Downarrow^{HH'} b$$

as noted by Denning [75], this property does not always hold, and can indicate that information may be transmitted even when no such transmission is possible. In section 9 we will explore just such a situation and develop a technique suitable for handling it correctly.

Section 8. AN EXAMPLE OF STRONG DEPENDENCY INDUCTION

This section illustrates Strong Dependency Induction by proving the correctness of a solution to the Confinement Problem in a simple system. Imagine a system where each object contains two disjoint subcomponents, one holding "data" and one holding a pointer to another object: The system has two sets of operations:

$\delta_1(y,x)$: if $y.ptr = x$ then $y.data \leftarrow x.data$
 $\delta_2(y,x)$: if $y.ptr = x$ then $y.ptr \leftarrow x.ptr$

If y points to x , then execution of $\delta_1(y,x)$ will copy data from x to y . If y points to x and x points to w , then after execution of $\delta_2(y,x)$, y will point to w as illustrated in figure 2.




Figure 2

If it is assumed that any object that is not "private" is a "spy", the Confinement problem can be stated as:

$$(\forall x,y) (x \Downarrow_{\varphi} y \supset \text{Private}(x) \supset \text{Private}(y))$$

I'll show that as long as a spy doesn't point to a private object, confinement is enforced. Formally the initial constraint φ can be chosen to be

$$\varphi(\sigma) \equiv (\forall y) (\text{Private}(\sigma.y.ptr) \supset \text{Private}(y))$$

That is, if in any state σ , an object y points to a private object, then y must be private

φ is autonomous and can easily be shown to be invariant (the only interesting case is $\delta_2(y,x)$).

Next, pick

$$q(x,y) \equiv \text{Private}(x) \supset \text{Private}(y)$$

noting that q is both reflexive and transitive. It can be easily shown that

$$(\forall x,y,\delta) (x \Downarrow_{\varphi}^{\delta} y \supset q(x,y))$$

By Lattice Induction, confinement is enforced.

Section 9. SEPARATION OF VARIETY

Strong Dependency Induction may fail to be useful in proving the absence of information transmission if the Strong Dependency relation is not transitive. This section explores two non-transitive

examples and derives a new technique, Separation of Variety, which extends the applicability of Strong Dependency Induction.

Consider the system

δ_1 : if q then $m \leftarrow a$
 δ_2 : if $\neg q$ then $b \leftarrow m$

Information cannot be transmitted from a to b in this system. If q is false, no information can be obtained from a ; if q is true, no information can be stored in b . Formally, it can be shown that

$$\neg a \Downarrow^{\delta_1 \delta_2} b \text{ even though}$$

$$a \Downarrow^{\delta_1} m \text{ and } m \Downarrow^{\delta_2} b$$

Thus an attempt to prove $\neg a \Downarrow b$ directly by End point Isolation is doomed to failure, for

property 3 requires that

$$(\forall \delta, x \neq a) (\neg a \Downarrow^{\delta_1} x) \quad [\text{yet } a \Downarrow^{\delta_1} m] \text{ or}$$

$$(\forall \delta, x \neq b) (\neg x \Downarrow^{\delta_2} b) \quad [\text{yet } m \Downarrow^{\delta_2} b]$$

The informal explanation above which argued that information could not be transmitted from a to b proceeded by distinguishing two cases - q false and q true. This can be formalised in the following way. Let

$$\varphi_1(\sigma) \equiv \neg \sigma.q$$

$$\varphi_2(\sigma) \equiv \sigma.q$$

Though $a \Downarrow^{\delta_1} m$, $\neg a \Downarrow_{\varphi_1}^{\delta_1} m$ and

$$(\forall \delta, x \neq b) (\neg a \Downarrow_{\varphi_1}^{\delta} x)$$

so by property 3, $\neg a \Downarrow_{\varphi_1} b$.

Similarly $\neg m \Downarrow_{\varphi_2}^{\delta_2} b$ which is the main result needed to demonstrate $\neg a \Downarrow_{\varphi_2} b$.

Since φ_1 and φ_2 cover all cases, $\neg a \Downarrow b$.

Not every separation into cases is legitimate. Consider the system

δ : if a then $b \leftarrow 0$ else $b \leftarrow 1$

Pick

$$\varphi_1(\sigma) \equiv \neg \sigma.a$$

$$\varphi_2(\sigma) \equiv \sigma.a$$

Then $\neg a \Downarrow_{\varphi_1} b$ and $\neg a \Downarrow_{\varphi_2} b$ yet information certainly can be transmitted from a to b . The trouble is that φ_1 and φ_2 are dependent upon the value of a - the very object that is in question as an information source. Separation of variety is legitimate only when each of the separating constraints is independent of the information source in question. Define

φ is a -independent \equiv_{def}

$$(\forall \sigma_1, \sigma_2) (\sigma_1 \stackrel{a}{=} \sigma_2 \supset \varphi(\sigma_1) = \varphi(\sigma_2))$$

If there is any set of constraints $\{\phi_i\}$, each a-independent, not necessarily disjoint, but such that any state is satisfied by at least one ϕ_i - then - if no information is transmitted from a to b given each of the ϕ_i 's in turn, then no information can be transmitted from a to b at all.

PROPERTY 5 (Separation of variety)

If $(\forall \sigma) (\phi_i(\sigma))$ [to guarantee that ϕ_i covers all cases]

and $(\forall i) (\phi_i \text{ is a-independent})$

then $(\forall i) (\neg a \xrightarrow{\phi_i^H} b) \supset \neg a \xrightarrow{H} b$

or more generally

PROPERTY 6 (Separation with Constraint)

If $(\forall \sigma) (\phi_i(\sigma))$ [to guarantee that ϕ_i covers all cases]

and $(\forall i) (\phi_i \text{ is a-independent})$

then $(\forall i) (\neg a \xrightarrow{\phi_i^H} b) \supset \neg a \xrightarrow{\phi^H} b$

Another use of separation of variety arises in showing that no information can be transmitted from a to b in the system

$$\delta_1: m \leftarrow m + k*a$$

$$\delta_2: b \leftarrow m \text{ mod } k$$

Pick $\phi_i(\sigma) \equiv (\sigma \cdot m \text{ mod } k) = i$

First $(\forall i) (\neg m \xrightarrow{\phi_i^{\delta_2}} b)$ though $m \xrightarrow{\delta_2} b$

more generally $(\forall i) (\forall x \neq b, \delta) (\neg x \xrightarrow{\phi_i^{\delta}} b)$

Since each ϕ_i can be shown to be invariant, End point Isolation can be used to show that

$(\forall i) (\neg a \xrightarrow{\phi_i} b)$

Finally, by Separation of variety $\neg a \xrightarrow{H} b$.

This example is especially interesting, for showing $\neg m \xrightarrow{\phi_i^{\delta_2}} b$ takes advantage of the fact that information transmission may be eliminated by sufficiently reducing the variety in the source (section 5).

Section 10. CONCLUSION

This paper has presented Strong Dependency, a formalism for describing information transmission in computational systems and for proving that information is not transmitted over certain paths. Attention was restricted to systems constrained by a class of predicates which are both autonomous and invariant and the paper considered whether any information at all might be transmitted over an information channel, without attempting to measure the channel capacity. A discussion of the more general situation may be found in [Cohen 76].

In analysing transmission from a to b, the paper considered the effect that a change in a might have on b over any arbitrary sequence of operations. If we are concerned with information transmission as the result of execution of a sequential program, we instead need consider the effect on b over executions of the program. The resulting theory is discussed in [Cohen 77].

ACKNOWLEDGEMENTS

It's a pleasure to thank Dorothy Denning for her comments and to thank the whole host of members of the CMU community for many discussions of this material.

REFERENCES

W. Ashby. "An Introduction to Cybernetics", 1956.

E. Cohen. "Strong Dependency: A Formalism for Describing Information Transmission in Computational Systems", Carnegie-Mellon Univ., Comp. Sci. Tech. Report, August 1976.

E. Cohen. "Information Transmission in Sequential Programs", Submitted for publication 1977.

E. Cohen, D. Jefferson. "Protection in the HYDRA Operating System", Proc. 5th Symp. on Operating Systems Principles, November 1975.

D. Denning, "Secure Information Flow in Computer Systems", Ph.D. Thesis, Comp. Sci. Dept., Purdue Univ., May 1975.

D. Denning, "A Lattice Model of Secure Information Flow", CACM v. 19, 5, (May 1976).

D. Denning, P. Denning, "Certification of Programs for Secure Information Flow", Purdue Univ., CSD-TR 181, March 1976.

A. Jones, R. Lipton. "The Enforcement of Security Policies for Computations", Proc. 5th Symp. on Operating Systems Principles, November 1975.

B. Lampson. "A Note on the Confinement Problem", CACM v. 16, 10 (October 1973).

J. Millen. "Security Kernel Validation in Practice", CACM v. 19, 5 (May 1976).

L. Rotenberg. "Making Computers Keep Secrets", Ph.D. Thesis MIT, MAC-TR-116, September 1973.

K. Walters, et al. "Structural Specification of a Security Kernel", 1975 Intl. Conf. on Reliable Software, April 1975.

APPENDIX: NON-AUTONOMOUS CONSTRAINTS

Strong Dependency is not a wholly adequate formalism for determining information transmission given non-autonomous constraints. Consider the system

$\delta: b \leftarrow a1$

constrained by

$\varphi(\sigma) \equiv \sigma.a1 = \sigma.a2$

It can be shown that $\neg a1 \not\triangleright_{\varphi} b$. However, information is certainly transmitted from $a1$ to b . An observer of b , knowing that φ holds, still gains information about $a1$ (its value!) from execution of δ .

The definition of Strong Dependency considers variety in $a1$ alone, but φ eliminates this variety by forcing $a1$ to remain the same as $a2$. The effect is similar to constraining $a1$ to be a constant. φ essentially spreads $a1$'s variety to $a2$. To deal with the spread, Strong Dependency can be extended to allow a set of objects as an information source, and autonomy can be considered relative to that set alone [Cohen 76]. In the example above, φ is said to be $\{a1, a2\}$ -autonomous and $\{a1, a2\} \not\triangleright_{\varphi} b$.

Formal methods for showing that b depends on $a1$ - without needing to explicitly be concerned about $a2$ - can be derived using statistical inference methods. That is, however, beyond the scope of this paper.

