Approved for Public Release; Distribution Unlimited
Case #06-0241

The Velocity of Information

Michael L. Kuras
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730
mlk@mitre.org

Salim K. Semy
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730
ssemy@mitre.org

Scott L. Surer
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730
surer@mitre.org

Abstract **-** *How can the health (or the robustness) of a complex information processing system be quantitatively scored? This paper proposes an answer to that question. That answer is a hypothesis that must be tested. A proposal for doing so is provided. The motivation for and the reasoning leading to this hypothesis is also summarized. This is an abridged version of a paper that will be available on the MITRE website.*

Keywords: complex-system, health, robustness, velocity of money, velocity of information, economy, economies, information technology, air operations center, time sensitive target, characteristic period

1 Introduction

Increasingly, Information Technology (IT) based automation is supporting military decision making. This automation and its integration into decision making are both seen as being increasingly complex. A case in point involves the prosecution of a class of targets known as Time Sensitive Targets (TST). Such targets require extremely rapid prosecution because they pose (or will soon pose) a disproportionate threat to friendly forces and objectives, and because their existence or value as targets is fleeting. Good decision making, rather than weapon lethality, is generally the key to successfully prosecuting such targets. IT based automation is now essential in such decision making and improvements to this automation are being constantly sought to consequently improve decision making. What can be done to quantitatively gauge the advantages of changes to this automation?

There are many facets to decision making and to the automation that supports decision making. One of these facets might be labeled effectiveness. How effective is the decision making, and how effective is the automation in supporting that decision making? Answering such questions must take into account what the decision making is about (or what the automation is supposed to do). This is ultimately a question of purpose and cannot be answered without explicitly accounting for this purpose. It is case specific.

Another facet is centered on how healthy or robust the decision making (or the automation supporting it) is. This facet can be examined in many ways that are entirely independent of the purpose of the decision making (or of the purpose of the automation supporting that decision making). Consider a situation involving an F-22 pilot. The pilot needs to make many split-second decisions. Although there are many factors involved, one of them involves the reflexes of the pilot. There are now many medical techniques for quantifiably gauging the reflexes of any human being, including pilots. Within limits, if pilot reflexes can be speeded up, the decision making of the pilot will be improved, regardless of the specific decisions being made. The pilot's inherent decision making capability is understood to be healthier or more robust.

How can an assessment of the health of IT based automation be made without recourse to purpose? This question applies to both subjective and quantified assessments of health. Without inquiring here into how such subjective assessments might be made, it can still be agreed that such assessments are made – and made over and over again. Given that such subjective assessments are now routine, can such assessments also be made quantitative in some fashion? Specifically, can they be made quantitative for the increasingly sophisticated IT based automation now being employed by the DoD? Even more specifically, if substantive changes are made to such automation, have those changes made the automation more or less robust or healthy?

2 The Hypothesis

This section summarizes points drawn from an analogy between economies and complex information processing systems (CIPS). Economies and CIPS are both examples of social complex-systems. All complex-systems exhibit interesting functionality at more than one scale. [ref. 1] For economies, the two most frequently distinguished scales are termed micro and macro (the micro economic scale and the macro economic scale of the same economy). Such a distinction is seldom made at present for a CIPS. Nonetheless, it is hypothesized here that an equivalent situation obtains for a CIPS. This paper offers a hypothesis centered on the **velocity of information** in a CIPS at what might be called its macro scale. Stated briefly the hypothesis is this: as the velocity of information increases for a CIPS, the CIPS becomes healthier or more robust. The analog to the velocity of information in an economy is established first, and then a way to compute values for the velocity of information is proposed, based on the treatment of its analog in an economy.

More fully stated the hypothesis is this. Within limits,[1] the greater the velocity of information (**V_I**) in a CIPS, the healthier it is. **V_I** can be quantified and approximated by taking the ratio of the quantity of data presented to the human operators of a CIPS and the number of invocations of objects in the autonomic portion of the CIPS.

In order to understand the hypothesis, the underlined terms need to be elaborated.

Complex Information Processing System: An information processing system that includes computers, power, connectivity equipment, software, and especially people in non-incidental ways. A complex information processing system is a social complex system in which information technology is a significant component. For the discussion that follows, the CIPS is divided into a human and an autonomic portion.

Health: an overall assessment of the ability to function, or the absence of impairments to functionality, most generally without regard to specific functions or purpose.

2.1 Origins of the Hypothesis

As noted above, economies are social complex-systems that function at multiple scales. Economists call one of these the macro economic scale; and one of the measures that they use at this scale is called by them the **velocity of money** (VoM). [ref. 2] As a concept, the velocity of money has developed and matured over the last 100 years and now serves as an overall gauge of changes in the health of an economy. The VoM for an economy is defined as the rate[2] at which the aggregate money supply in an economy circulates or flows as viewed at the macro economic scale.

By way of illustration, the VoM for the US economy has changed from roughly 7 times per year to more than 9 times per year in the period from 1980 to the present. The US economy is substantially healthier today than it was in 1980.

[1] These limits are not examined in this short paper.

[2] The use of the word velocity to refer to a rate will be awkward for many engineers. Regardless, it is important to not impute many of the familiar connotations associated with the word velocity to the notion of the velocity of money as meant by economists. The velocity of money is a rate or a frequency of usage; it is not any sort of "distance" divided by "time." It is also imprudent to assume that the rate involved is "regular," for example. It may be, but that is not automatically implied.

The VoM is a macro economic measure. It is a measure that must be interpreted at one scale of an economy (the macro economic scale) even though the data on which it is based is collected at another scale (the micro economic scale) of the same economy.

Economies are social systems that function at multiple scales. Although the functionality at different scales is related, there is no fully reductionist and constructionist way to account for or to fully characterize this interdependence. In particular, the rate at which a particular dollar is used and reused within an economy, or even that of all individual dollars, does not fully account for the VoM and its relationship to the overall health of an economy. The interpretation of VoM applies only to the aggregate of money at the macro economic scale and not to what is perceived as many individual units of money at the micro economic scale and the many individual economic exchanges in which that money is used. A unit of money is visible at the micro economic scale; it is not visible at the macro economic scale. The unit of the aggregate money supply is visible at the macro economic scale; it is not visible at the micro economic scale. Moreover, the notion of health or robustness that can be applied to the whole economy at its macro economic scale cannot be subdivided and associated in any fractional way with the functionality visible at the micro economic scale.

In an economy, the two most important characteristics are the movement and the accumulation of value. In the case of a CIPS what are important are the movement and the accumulation of information. VoM is closely related to the rate at which value moves within an economy. VoM serves as a convenient and accessible means of approximating what might be termed the **Velocity of Value** (**V_V**) in a monetary economy. In the case of a CIPS, the exact analog of the Velocity of Value is the Velocity of Information **(V_I)**. The next section outlines how the V_I for a CIPS might be determined. This method is the analog of how VoM is computed and used as an engineering approximation of the Velocity of Value in an economy. It is the Velocity of Value at the macro economic scale that is the true gauge of the health of an economy, and not VoM. The VoM serves as an engineering approximation of the Velocity of Value.

The VoM in an economy is computed as AFP / M. AFP is Aggregated Finished Product and M is the money supply. Economists actually use many nuanced variants of AFP (such as the gross national product, the net domestic product, and so on). The same is true for M. In this case, the nuanced variants are termed M_1, M_2 and so on. The use of AFP and M here is done to avoid any implications that might be associated with any of the nuanced variants actually used by economists.

AFP is the aggregation of all finished goods and services generated within an economy during a nominal or standard interval (such as a year) and that is expressed in terms of the prices paid for them. Prices are scored in units of a currency such as the dollar. M is also scored in units of a currency. It is all of the currency in circulation in an economy. It is usually treated as a constant.[3] M is basically all of the money in our collective wallets but treated collectively only.

The qualification of "finished" in AFP is used to focus the scoring on just the goods and services that are actually delivered to final (ultimate) consumers in the nominal interval, AND that originated with their initial producers in the same interval of time. This is done to preclude multiple scorings of the same goods and services that are exchanged multiple times as they pass through a chain of producers and consumers while being augmented in the process. Any additional goods or services that are attached along the way are also counted (but just once) in the scoring of AFP.

The foregoing notes the most basic elements[4] of how economists quantify and interpret a characteristic of a whole economy (its health). Selected data is gathered at the micro economic scale and then interpreted at the macro economic scale. This interpretation applies to the whole economy. It is visible at the macro economic scale but it is invisible, in whole or in parts, at the micro economic scale at which the data on which it is based is collected. The next subsection outlines a method to gauge the health of a whole CIPS by making analogous measurements and interpretations at two distinct scales of the CIPS.

2.2 Measuring V_I for a CIPS

What is proposed here is that $\mathbf{V_I}$ can be quantified using the following equation.

$$\mathbf{V_I \approx V_\theta = AFI / (\theta\ supply)}$$

The Velocity of Information ($\mathbf{V_I}$) can be approximated by V_θ, the rate at which the aggregate of actual computational activity circulates or flows through a CIPS. V_θ is the close analog of VoM. V_θ can be computed as the Aggregated Finished Information (AFI), the aggregate of information bearing data presented to the human operators in a CIPS, divided by the aggregate computational activity within the CIPS (θ supply).

In a monetary economy and at the micro economic scale, an individual economic exchange involves a producer and a consumer swapping a good or service for a quantity of money. In a CIPS, money and goods are not exchanged. Instead, objects in a CIPS can be viewed as exchanging computational activity for strings of bits. Objects can be individual people or pieces of the autonomic portion of the CIPS consistent with the Object Oriented Programming System (OOPS) paradigm. Information is associated with the ordering (or the order) in a bit string. It is not the bit string itself. Higher dimensional data objects can be converted into bit strings. A random bit string carries no information. The longer an ordered string, the more information it carries.

The AFI is the analog of the AFP in an economy, and it can be scored by counting all of the ordered bits that are delivered to the people of the CIPS during this period. AFI can be expressed in units such as Gigabits/hour.

To exactly match the economic model, it would also be necessary to include what might be termed the "information added" contribution to the AFI that is the analog of the so-called "value added" contribution to the AFP in the production chains within an economy. It is assumed here that this contribution to the AFI is made in the autonomic portion of the CIPS but is dominated by the contributions from the final or end human operator contributions to the AFI. This is basically because most of the meaning or information content in the data does not actually materialize until presented to human users. This is due primarily to the extraordinary context provided by trained human operators.

Scoring the θ supply has to account for the contribution from both the autonomic and human portions of the CIPS as well. Although the human contribution is important, it is assumed here that human contribution to the actual θ supply is dominated by the contribution from the autonomic part and can be ignored.

This may seem at first glance to be contradictory to the previous simplification concerning the relative contributions to AFI of the human and autonomic portions of the CIPS, but it is not. In a contemporary CIPS, most or all of the computational activity that is in the autonomic portion of the CIPS is of a different kind than in the human portion (it lacks the extraordinary context provided by human operators, for example[5]), but there is much more of it. On the other hand, most of the meaning (information content) that might be attached to bit strings does not materialize until actually presented to human operators.

For the autonomic contribution, floating point operations (FLOPS) can be used to quantify computation. The choice of another unit is possible as well. Regardless, it is important that what is frequently termed "idle cycles" when quantifying computational activity not be included. FLOPS can be understood as being used to generate the

[3] The money supply isn't constant. Manipulating it has been a frequent means of engineering an economy.

[4] Many other important elements, such as the treatment of inventories, have been omitted in this short paper.

[5] In fact, the respective computations occur at different scales.

information (as strings of ordered bits) that is eventually presented to human operators.

At the micro economic scale, an individual economic exchange can be represented by the equation

$$m \approx x_i (g)$$

where m is expressed in terms of a currency such as dollars and g is the quantity of a good or service specific to the exchange, denoted by i. x_i is the function that establishes the equivalence of the specific exchange. This equivalence does not have to be an equality.

It is not necessary or appropriate to assume that the money involved in any specific economic exchange captures some absolute measure of value. Nor is it necessary to assume that the value in the money and the value in the goods or services are the same, or that the amount of money is repeated exactly in other exchanges involving the same goods and services or participants. The producer and the consumer negotiate an exchange, and they execute the exchange when both parties are satisfied with the equivalence. Actual negotiations and executions can be extremely perfunctory and automatic as well as protracted and deliberate. This equivalence is local or specific to just this exchange, and is not global in any sense.

At what might be termed the micro scale of a CIPS, it is possible to imagine two objects in what can be termed the analog of a micro economic scale exchange. In this case, however, one object "receives" FLOPS and "gives" the other object a string of ordered bits in return.

At this micro scale in a CIPS, the following equation holds:

$$F \approx e_k (B)$$

where F would be expressed in FLOPS and B in ordered bits, and e_k is specific to the equivalent of a micro economic exchange, and k ranges over all such exchanges. At the macro scale all such transactions can be aggregated as a statistical function, **E**, and applied to the aggregation of all ordered bits and FLOPS exchanged:

$$\Sigma F = \mathbf{E} (\Sigma B)$$

There is no reason to assume that **E** does not change with time. However, for a number of reasons, including the immediate interest in a self-relative measure (i.e., is a CIPS getting healthier or not; and not, is it healthier than another one), the nature of the relationship between ΣF and ΣB (i.e., the statistical function **E**) can be assumed to change by no more than some linear adjustment. As a result, a FLOP can be treated as the equivalent to an ordered bit except for some scalar adjustment.[6] In other words, both **E** and e_k can be understood as functions that map one unit (ordered bits) to another (FLOPS).

This also suggests that invocations (as distinct from the extent or intensity of any computations following invocations) are what really matter. Invocations can be counted instead of discrete computational operations in scoring the θ supply. This observation could also be extended to the people in a CIPS (which have been deliberately ignored in this discussion of scoring the θ supply). Their contribution (assumed to be negligible here) could be reintroduced by counting the times each person "does" something.[7]

The θ supply has to be treated as a constant (whether inclusive or exclusive of human activities in the information system). This in effect determines the duration in which the counting of ordered bits is performed in scoring AFI. This then is the recipe suggested:

1. Choose a convenient definition of an object that permits the counting of all of the invocations of the set of objects {objects} in the autonomic portion of a CIPS.

2. Measure the duration (T) that elapses to accumulate a fixed number of invocations (N) of the autonomic objects in the CIPS. Also count the number of ordered bits (B) that are presented to the human operators in the information system in this duration T.

3. $\mathbf{V_\theta = (B / T) / N}$

2.3 Generalizing the Hypothesis

The foregoing is based on an analogy between two different kinds of social complex-systems, economies and CIPS. In the case of an economy, the VoM is used as an engineering approximation of $\mathbf{V_V}$ and can be used as a

[6] Recognition of such a statistical function is not unique to a CIPS. It should be apparent in other complex-systems such as economies. There is no reason to assume, however, that such a statistical function is generic. It could be specific to types of complex-systems, or even to specific instances of complex-systems. It could be a signature function.

[7] More technically, what would be counted is the number of times that a person executes his or her own personal OODA-loop. OODA-loops are the John Boyd template for analyzing the decision making process: observe, orient, decide, and act. These are not steps; rather they are overlapping regimes in the overall activity of an autonomous decision making agent at some scale of a more extensive system.

gauge of changes in the health of an economy. In the case of a CIPS, $\mathbf{V_\theta}$ is an engineering approximation of $\mathbf{V_I}$ and can be used as gauge of the health of a CIPS.

The premise underpinning the hypothesis discussed here extends to all complex-systems. This can be emphasized by considering the inverse of measures such as these, and labeling them the **Characteristic Period** (**P**) of their respective systems. In the case of an economy, P is 1 / $\mathbf{V_V}$, and in the case of a CIPS, it is 1 / $\mathbf{V_I}$. And P can be approximated in these particular cases by using VoM and $\mathbf{V_\theta}$ respectively.

Doing this deemphasizes the specific measurements that must be taken to calculate the Characteristic Period of a complex-system while emphasizing an inherent similarity shared by all complex-systems – that their behavior or functionality at their dominant scale is periodic while operating at or near equilibrium, and that the health or robustness of such systems improves as this natural periodicity shortens.

3 Testing the Hypothesis

This hypothesis can be tested using a portion of the Air and Space Operations Center (AOC [3]) CIPS. An AOC is a mature element in the overall (social) system employed by the US to conduct military operations. The AOC CIPS supports the planning, executing, and assessing of the theatre-wide air and space operations performed by the AOC. Approximately 3,000 personnel, 70,000 sq. ft. of space, and $50 million of computer and network automation are needed to operate an AOC. The AOC operates on a daily rhythm, producing and executing a daily Air Tasking Order (ATO).

The AOC also plays a role in the prosecution of Time Sensitive Targets (TSTs). The TST process is a microcosm of the daily process, involving many of the same planning, executing and assessment functions but in a compressed timeframe in order to prosecute TSTs. A readily identifiable fraction of the overall AOC is devoted to this role. It is proposed to instrument the TST portion of an AOC CIPS in order to test the hypothesis because it is well understood and is a more manageable CIPS to study than the entire AOC CIPS.

A quantitative health assessment (such as the one proposed here) should align with the qualitative or subjective assessments made by subject matter experts (SMEs) familiar with the AOC and TST prosecution. Making such comparisons, over a prolonged period, is the essence of how we would propose that the hypothesis be tested. If the hypothesis is correct, it will serve to refine and elaborate what is already the subjective assessment of a CIPS' health.

For the TST CIPS, the number of ordered bits (B) could be measured in a number of ways, such as accumulating the volume of data written to the display interfaces of the consoles used by TST personnel. The number of object invocations (N) could be measured by accumulating the number of application process activations initiated by the operating systems of the computers writing to the display interfaces. Over an extended period of time, the ratio of (B/T)/N could be compared with a subject matter expert's subjective opinion of the health of the TST CIPS, to reveal any correlation between the quantitative metric and the qualitatively assessed health.

4 Conclusions

This paper presents a hypothesis: that the self-relative health of a CIPS can be gauged by monitoring changes in its Velocity of Information, or equivalently in its Characteristic Period. This is based on a strong analogy between an economy and a CIPS. This is not to suggest, however, that everything about an economy also applies to a CIPS. It is important to understand where the analogy does hold and where it does not in order to better understand how to measure the health of complex-systems in general. The hypothesis of this paper should be tested. An approach has been outlined to do this using the Time Sensitive Targeting Cell of an Air Operations Center.

Throughout the paper, there were also several ideas that were noted in passing. For example, is there a statistically meaningful pattern regarding the distribution of computations and the quantity of ordered bits that objects in the autonomic portion of a CIPS generate? Is such a pattern generic, or is it a signature of each CIPS? How can one quantify human thinking? Is the Characteristic Period a generic characteristic of all complex-systems? These questions should be revisited to more thoroughly assess this hypothesis and the many issues related to it.

Finally, improvements to a CIPS can directly impact the performance of the decision making process that the CIPS supports. Validating the hypothesis of this paper may have broader implications, by pointing to ways to gauge improvements in those more inclusive processes themselves. In the case discussed, the health of the CIPS in the TST cell does not automatically imply the health of the TST cell itself but may still be an important factor. Nonetheless, understanding the relationship between the characteristic period of a portion of the AOC and its CIPS might well provide insights into how to gauge and even improve the overall AOC process or of portions of that process.

References

[1] Kuras, M.L., and Brian E. White, Engineering Enterprises Using Complex-System Engineering, INCOSE, 2005.

[2] Samuelson, Paul A; Economics; multiple editions; Mcgraw-Hill.

[3] Air Force Instruction 13-1AOC, Volume 3, August, 1 2005, http://www.e-publishing.af.mil/pubfiles/af/13/afi13-1aocv3/afi13-1aocv3.pdf