

VALUED INFORMATION AT THE RIGHT TIME (VIRT): WHY LESS VOLUME IS MORE VALUE IN HASTILY FORMED NETWORKS

Frederick Hayes-Roth¹
Naval Postgraduate School
Monterey, California

ABSTRACT

Defense and homeland security leaders have focused recently on the problems of fielding networks to enable rapid decision-making and agile responses to various crises. Mostly they have concentrated on the lowest levels of networks, namely the hardware and software to enable bits to flow from senders to receivers. However, most crises require a different approach, one that emphasizes the highest levels of network design. At these levels, the problems we focus on are: **Who** needs **What** information, and **How** does that information **Find** them? In addition, because people in crises have so little time, we must also answer this question: **How do we assure receivers are not glutted** by a deluge of low-value data and consumed by attendant low-value tasks? Our answers to these questions employ dynamic context and operator requirements to assure that high-value information flows quickly where it's needed and is processed promptly by recipients. We call this approach Valued Information at the Right Time (VIRT). Initial studies have shown that this approach reduces the volume of bits by several orders of magnitude. It also raises the productivity of every operator enormously by assuring each can give immediate attention to truly valued information. A VIRT perspective leads us to see networks as information supply chains. Well-designed supply chains will dramatically improve the performance of hastily formed networks (HFNs).

WHAT & HOW OF NETWORK EFFECTIVENESS

This paper addresses the question of how to *implement* information superiority, especially in networks of organizations, people and machines that have limited resources. That is, can we make our military operators and homeland security agencies *superior* to all competitors and against all challenges by providing them better information? What information would that be? What kind of process would deliver that information? What kind of process would cause the operators to act promptly on the best information? How would we avoid likely pitfalls, such as glutting operators with extraneous bits and assure

they spent their time mostly acting upon important and timely ones? These basic questions weren't addressed adequately in the early formulations.

VADM Cebrowski was an early proponent of network-centric approaches to warfare and operations other than war. His concepts are embraced in DoD visions for using information superiority as a foundation for better, faster, more effective military operations. The "bible" for these concepts is the insightful book by Alberts, Gartska and Stein [1]. As an introduction to this paper, I'd like to give a highly simplified description of **what** information superiority is and **how** such leaders suggest it will be achieved (cf. [2]). Responsibility for implementation of these ideas has, of course, passed from the seminal thinkers to such organizations as OSD's NII, the Navy's FORCEnet, DISA's GIG/NCES, and DHS advanced technology groups. While people may differ on details, the basic ideas follow:

Networking makes it possible to communicate quickly, across great distances, and among diverse services and agencies. Thus, it should be possible for everybody to access all relevant information, regardless of who produced it or where.

Information superiority would then result from each operator finding and accessing all information relevant to the mission. To expedite this process, each supplier should provide meta-data describing the contents and qualities of the data supplied.

Bringing this down to a technical implementation level, all data would be marked up with some XML tags reflecting a supplier's categories for content and quality. An operator seeking information would describe the data properties desired. A network query would search for matching meta-data and then retrieve the data responsive to the operator's query. Thus, the network would deliver all relevant data, with minimal delay, overcoming traditional barriers to effective information sharing.

Several additional qualities would emerge from this effective dissemination of data. All participants in an operation could share a common operational picture, assuring that all would sense and respond to the same

¹ The author is a Professor in the Information Sciences Department at NPS. His email is: hayes-roth@nps.edu. This work has been supported in part by the Navy Cooperative Engagement Capability (CEC), PEO Integrated Warfare Systems 6, and the Naval Postgraduate School.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2006		2. REPORT TYPE		3. DATES COVERED 00-00-2006 to 00-00-2006	
4. TITLE AND SUBTITLE Valued Information at the Right Time (VIRT): Why Less Volume is More Value in Hastily Formed Networks				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School ,Monterey,CA,93943				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Defense and homeland security leaders have focused recently on the problems of fielding networks to enable rapid decision-making and agile responses to various crises. Mostly they have concentrated on the lowest levels of networks, namely the hardware and software to enable bits to flow from senders to receivers. However, most crises require a different approach, one that emphasizes the highest levels of network design. At these levels, the problems we focus on are: Who needs What information, and How does that information Find them? In addition, because people in crises have so little time, we must also answer this question: How do we assure receivers are not glutted by a deluge of low-value data and consumed by attendant low-value tasks? Our answers to these questions employ dynamic context and operator requirements to assure that high-value information flows quickly where it's needed and is processed promptly by recipients. We call this approach Valued Information at the Right Time (VIRT). Initial studies have shown that this approach reduces the volume of bits by several orders of magnitude. It also raises the productivity of every operator enormously by assuring each can give immediate attention to truly valued information. A VIRT perspective leads us to see networks as information supply chains. Well-designed supply chains will dramatically improve the performance of hastily formed networks (HFNs).					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 10	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

perceived reality. Because they could all “sing from the same sheet of music,” they could self-synchronize, enabling higher levels of autonomy and agility.

So the catechism goes more or less like this: *What is information superiority?* A state where each operator acquires all relevant information in a timely way. *How is information superiority achieved?* Enabling each operator to access quickly all relevant information leads directly to shared awareness, better decisions, and greater agility.

Unfortunately, we find much of this argument naïve and likely harmful. As a high-level motivational concept, it has served its purpose well. As a guide to implementation, it’s seriously misguided. The purpose of this paper is to show why that is the case and how to achieve orders of magnitude better communication with much higher operator productivity. Such improvements, in turn, should lead to much higher operational effectiveness. In contexts where organizations must cooperate and align in ad hoc crises, the challenges are even greater, so the rewards for efficient communication should be even greater. In the next sections we revisit the basic tenets of information superiority, diagnose the naïve and dangerous misconceptions, and then proceed to recommend remedies.

SHARED AWARENESS & RAPID DECISIONS

One aim of information superiority is to support a shared and common situational awareness among collaborating forces. The Star Trek creature, the Borg, in some ways represents the ideal here: every member of the Borg participates in a common distributed cognition. When the Borg assimilates new units, they too become part of the group mind. In command-control parlance, all entities read and write from one Common Operational Picture (COP).

Recently, people have begun to realize some of the fundamental reasons why such an ideal is not attainable. First, events occurring continuously at great distances can’t be instantaneously communicated. Second, all events are sensed and *interpreted* by various systems and people, and these interpretations are rarely 100% certain and accurate. Further, different perspectives produce multiple alternative interpretations. These can’t be instantly communicated, compared or resolved. Third, because the consumers of information are action oriented, each has customized needs based on their own context, concerns, and area of interest. At different levels of echelon, with different spatiotemporal spans of concern, decision makers literally need different types of information, reflecting appropriate levels of abstraction and aggregation. Across different military and domestic agencies, moreover, differences in concerns and viewpoints create demands for

varied types of information. No single encoding of information can address these disparate needs.[3] So people now realize each operator needs a tailored, i.e. *User Defined*, operating picture (UDOP).

Current efforts don’t yet address all of these concerns in a credible way[4]. Efforts to build a COP and display to each operator an appropriate UDOP don’t yet squarely face the problems of uncertainty, alternative interpretations, latency, or the need for different levels and different types of information. There’s a tacit belief that these problems are marginal and that the best overall approach is to move, fuse, and disseminate as much information as possible. Following this logic leads to the belief that if we make all information quickly available, good things will follow. Specifically, when each person has timely access to all relevant operational information, rapid and effective decision-making will ensue. As an important side-effect of such assumed optimal decision-making, diverse participants will self-synchronize by referring to their shared COP.

I find this idealized concept implausible, because it doesn’t address any of the fundamental limitations faced by distributed, diverse entities trying to get their own important work done well. Having studied decision-making in small and large organizations, in normal times and in crises, I think we must address these limitations realistically to improve actual performance. Specifically, we need to appreciate that humans are operating at finite information processing speeds, far below those required to achieve Borg-like total shared awareness. Moreover, because of all available information only a tiny fraction is actually valuable for operators, it’s foolish to focus on making them process it all. We ought to ask how we can assure they apply their scarce mental resources to high-value information. Two concepts that underlie the answer are MCNs and VIRT, the subjects of the next two sections.

MODEL-BASED COMMUNICATION NETWORKS

If we can’t possibly achieve Borg-like perfect distributed shared awareness because of fundamental limitations, what should be our practical objective? We should try to achieve three principal goals:

1. Each operator should expend most of his or her resources maintaining the best description of situational variables most important for the operator’s own task effectiveness.
2. Each operator should provide information to others that they would most value.
3. Communications should be optimized to maximize receipt and use of timely valuable information, as well as filtering out valueless information.

These capabilities require that communication become sensitive to the dynamic context of the receiver, because the receiving operator's current plans and beliefs determine what new data would be worth transmitting. When many operators work in the same environment, the potential to connect them with a network initially seems appealing. Naïvely, it seems attractive to have all data be disseminated to every interested party. There are many different network architectures for achieving this, but they all focus on moving all bits rapidly to interested parties.

When communication systems strive to continually update many receivers with dynamic situation data, the demands on bandwidth and processing can be enormous. Because each sensor or reconnaissance asset sends a steady stream of updates, every planner and operator finds a steady set of messages in the inbox needing attention. In real organizations, this produces a familiar sense of information glut [5], and operators regularly fall behind. Even in established organizations executing standard processes, backlogs of hours or days of "real-time" data are common. For HFNs rapidly assembled to respond to an ad hoc crisis, such delays translate directly into death, disease, and other significant lost opportunities. Obviously HFNs must respond rapidly. Thus we must not allow excessive traffic volumes to produce processing backlogs and delay timely decision-making.

We can improve the efficiency of communication resources by recognizing that receivers have different purposes and beliefs. What operators most need communicated to them is evidence that can affect those purposes or that should change those beliefs. For example, a pilot who has a planned route to a rescue site needs to know about adverse weather en route or newly corrected coordinates of the victims. Once new evidence has been conveyed, repetitions add no value. Evidence about irrelevant regions of space and time also lacks value.

Good communicators need to understand the dynamic context or "state" of the recipient. This state includes the recipient's mission, assumptions, and beliefs. When the sender understands this state and uses it to determine what to convey, communication is called "state-full" [6] Most network technology employs a state-less foundation, so that networks today are mostly ignorant of the receiver's context and what information the recipient would value.

Many operations of interest to us involve dynamic entities moving through space and executing activities that change over time. We often model these entities mentally, as when we dead-reckon an inferred position of a vehicle based on its previous reported location and velocity. When multiple

parties are operating in a distributed arena, they often wish to create and maintain a shared understanding of the "state" of the environment as well as enemy and friendly elements in that space. This gives an additional significance to state-full communication. If the beliefs of each party about such dynamic entities are reflected in computable dynamic models, each party should be able to update its overall situation autonomously by dead-reckoning its own models for the various entities.

Air traffic controllers, for example, maintain "the bubble" in their head: they continually project flight paths of all aircraft from current data into the future to assure no conflicts will arise. Fire officials maintain a model of winds, fuel loads, fire lines, evacuation centers, and lines of communication, for example. Each of these officials uses his or her own model of entities that are often controlled by and modeled by others. That is the nature of collaborating, distributed, cooperative communication. Models must be immediately available to support each decision-maker's specific needs for rapid projection of expected events critical to that decision-maker's concerns.

A Model-based Communication Network (MCN) is a state-full distributed system of collaborating nodes that maintains an optimal shared understanding of the situation [7]. The situation at each node is composed of models of all entities relevant to its mission. Each node can dynamically project future states of such modeled entities. Each node, in addition, understands the state of its collaborating nodes, including the others' missions, assumptions and beliefs that might be affected by changes in its own perception of the situation. Thus, node A is aware of what node B is trying to do and how node A's own knowledge might impact B's beliefs and behavior.

In principle, each node can determine what it knows that another node would want to know. That defines "valued information." In an MCN, we give priority to conveying valued information. We also try to eliminate transmissions of low-value information, because these consume valuable resources and increase latencies. Because communication in HFNs always stresses the limited resources, we need to make HFNs behave as much as possible as MCNs. While the concept of MCN is easily stated, turning it into operational practice is the goal of VIRT, addressed next.

VALUED INFORMATION AT THE RIGHT TIME

Although we can't know perfectly what every operator wants to know, military practice has developed standard procedures that we can build on. For example, commanders often dictate to their staffs their critical information requirements (CCIRs). Commanders consider

these conditions significant enough to wake them in the night, for example. Similarly, every operation that results from deliberate planning needs to know when assumptions underlying the plan no longer seem credible. A planned rescue flight, for example, can't go if weather forecasts now predict surprising sandstorms or thunderstorms. This basic idea is illustrated in Figure 1.




Figure 1. A simplified architecture for VIRT.

The architecture in Figure 1 simplifies much of the complexity by focusing on just a single plan, apparently planned and periodically adjusted by only the one person illustrated. Of course, in real organizations many teams pursue many objectives, so there are many planners and plans extant at any point in time. Nevertheless, the key elements of the VIRT approach appear in this simple view.

The overall flow in Figure 1 starts on the left side, where a planner has generated a plan. Each plan describes time-phased actions that should accomplish the plan's objectives. The planner considered what the state of the world would be at the time the plan executes, and the planner's beliefs about that future state correspond to the "assumptions" recorded in the plan. When planners select a plan, they usually evaluate it and compare its costs and benefits to other alternatives. They can record their reasons for selecting one particular alternative in the form of a *justification*. A justification usually explains how the plan accomplishes the objectives in a situation where the assumptions validly hold and, also, why the planners prefer the selected plan to the alternatives considered. The justification often reflects that all of the alternatives considered had greater risks or costs in comparison to the chosen one.

Let's consider a simple example. The planners might have an objective to rescue a small group of people on the ground in forested terrain. Their basic choices consist of reaching the people by ground or air and extricating them by ground or air. The likely options for ground transport include wheeled and tracked vehicles, horses, and humans. The likely options for air transport include rotorcraft v.

fixed wing aircraft. Given a number of factors, they quickly reject all but the following skeletal plan:

1. Survey the area by fixed wing aircraft to find the best landing spot for a helicopter.
2. Send a helicopter with a search and rescue (SAR) team.
3. Land the helicopter at the chosen site.
4. Find and recover the party using the SAR team.
5. Depart by helicopter and return the party to a chosen medical facility.

Given this skeletal plan, the planners then focus on possible aircraft and routes, total expected flight times and associated fuel requirements, and possible time sequences for the flights. The flights become highly dependent upon the assumed wind, visibility, and icing conditions en route and at the search and rescue area.

Let's complete the example plan. The planners assume that a 90-minute aerial survey will be required to choose the best landing site. They choose an available low-altitude aircraft that carries appropriate instruments and can reach the site in a two-hour flight. The aircraft has 6.5 hours of fuel, adequate for two 2-hour legs and a 1.5 hour survey, leaving enough fuel for a 1-hour safety reserve. Winds in the area are forecast to be excessive between the hours of 1300 and 1800 local, and adequate sunlight is expected only from 1000 to 1900. For these reasons the flight is planned for early tomorrow morning, so that the survey begins promptly at 1000. Thus, take-off is scheduled for 0800. The helicopter is scheduled for a 3-hour flight to the search area, and is planned to depart at 0900, so that it can receive landing coordinates at 1130 from the aircraft survey team 30 minutes prior to touching down.

Even this example leaves out countless details, but it provides enough to illustrate the key VIRT architecture features. The VIRT dependency monitor watches for changes in forecast or actual conditions that threaten the plan by undercutting its justification. In the current case, the dependency monitor in Fig. 1 needs to revalidate periodically the key assumptions regarding aircraft availability, aircraft capabilities, winds, visibility, icing and fuel consumption.

My USMC colleagues at NPS, under the leadership of LtCol Carl Oros [8], have shown how VIRT can be implemented in many standard operations. Basically, we begin by identifying assumed conditions that each phase of the operation depends on, such as healthy troops, effective weapons, suitable communications, and credible target location and identification. Each such assumption is negated to define a *condition of interest (COI)*, a type of worrisome event that warrants immediate notification. The

information network is tasked to monitor these COIs and to alert the operator immediately when one is detected. In short, the network becomes goal-driven, with the goal of helping each recipient receive the information it most values.

To implement VIRT, we need a vocabulary of terms, such as “coordinates of search area,” “aircraft operational status,” “flight crew readiness and availability,” “winds aloft,” “thunderstorm activity,” “sandstorm activity,” “expected time for coordinates availability” and “accuracy of location estimate.” These variables are associated with specific operations and units, and they may be indexed by space and time, because many of them are dynamic. A COI is then written as a kind of expression or “continuous query” that probes current information sources for a change from “non-event” to “event.” For example, when thunderstorms along the planned route arise, that’s a change of data value from uninteresting to important, because the COI just becomes true.

VIRT works by allowing operators to express their COIs in an easy-to-use language of expressions built upon their own vocabulary. COIs are monitored by brokers or agents that perform continuous queries on relevant data. Data processing can be pushed further upstream, close to relevant sensors, to optimize computation, minimize delay, and minimize bit flow. There are many opportunities to optimize the work expected on identifying and selectively transmitting the valuable information corresponding to the COIs. Regardless of how that work is done, VIRT focuses on assuring that such high value information moves through the network, with priority, to the affected and interested recipients.

In the end, each operator can delegate to a system with VIRT capabilities responsibility for monitoring all identified COIs. This reduces the operator’s workload by offloading routine monitoring of predictable problems. It reduces the flow of information to the operator. It frees up the operator’s attention to deal promptly with alerts. It also makes more operator time available to consider novel or unpredictable situations that may arise.

Shipboard electronic navigation systems illustrate VIRT principles. A navigator, based on his captain’s intent, informs the system of intended route, required safety margins, fuel consumption constraints, etc. The ship’s fathometer, radars, GPS receiver, compass, and engineering plant all monitor COIs and deliver an alarm when the navigator’s expectations are no longer true. A watch officer then intervenes appropriately. Merchant ships have exploited this superior information value chain by drastically reducing manning on their bridges. The next

section explains why VIRT improvements are much more than just incremental. In fact, they are extraordinary.

LESS DATA, ORDERS OF MAGNITUDE BETTER

Over the last two decades, huge improvements in manufacturing effectiveness were achieved through a combination of “just-in-time” deliveries and other “supply chain integration” techniques. The key behind these improvements was to make each process step as efficient as possible and to minimize idle resources. Optimal results could be achieved if expensive processors and people were always busy on high-value products and unfinished products moved smartly from processor to processor, eliminating idle time and inventory. In addition, if each stage packaged and delivered its results in just the way the next process step found most suitable, every step could execute quickly with minimal cost and delay.

It’s helpful to look at business and military decision-making systems as “information chain” integration tasks, analogous to supply chain integration. In information chains, we are moving *bits* rather than *molecules*, but we still have scarce and expensive resources that shouldn’t be wasted. In most military and emergency relief operations, our scarcest resources are time for decision making and communications bandwidth among mobile entities. We don’t want to waste these resources by moving low-value bits or by creating backlogs of unprocessed bits awaiting analysis. In these operational contexts the penalties for delay are magnified, because bits are “perishable.” Like fresh fruit, information is best consumed when ripe, before becoming stale. Furthermore, in military and relief contexts, decisions that are poorly informed or late often cost lives. Thus, the rewards for information chain efficiency will be even greater than for supply chain efficiency.

In a recent paper[9], I evaluated the quantitative advantage of a VIRT approach (“smart push”) for information dissemination and compared it to the best possible version of “smart pull,” where operators retrieve information from the GIG relevant to their missions. The scenario involved a helicopter pilot flying a mission in hostile territory (akin to that in [8]). Denning [4] paraphrased my analysis, and I reuse some of his pithy version in the paragraphs below.

Before starting, the pilot creates a flight plan that avoids storm cells and air defense positions. The pilot will deviate only on learning of changes in storm and defense positions, as well as movements of other aircraft, that intersect the planned flight path. Various other technologies (weather observation, radar) track storm movements, anti-aircraft positions, and other aircraft

through the entire region. Of all these data, however, the pilot will only value those that stimulate the pilot to consider promptly a deviation from the current route.

Consider a flight path through a region 200km on a side. Sensor resolution in the region is 1 km, giving 40,000 grid points. Vertically, data are available at 500m intervals from altitude 0 to 6 km, a total of 13 altitude coordinates. That gives 520K grid points in the 3-D volume. Forecasts of ten variables are tracked at each grid point, giving 5.2M data values in the volume. Weather forecasts are updated every 30 min., and the flight is scheduled for 4.5 hours, giving 10 update times. Thus the total size of the data space is approximately 52M values.

In a “dumb” push environment, the sensors and updaters send new information to the pilot whenever they get it, so during the 4.5 hr flight, the pilot receives all 52M values. If we instantiate a “smart pull” as described for example by Krieger on behalf of ASD NII [10], the pilot uses tools to search the data for items more obviously relevant to his interests. For example, he might discard data more than 5 km away from the flight path or set local filters to hide data that have changed less than 5% since their previous reading. Even if such filters remove 99% of the received values, the remaining 1% (520K potentially relevant values) will exceed the pilot’s capacity to make sense of them and constitute a likely distraction. Worse, the 99% of values discarded wasted scarce bandwidth and probably slowed deliveries to other warfighters.

In a smart push environment, the pilot describes COIs so that data outside some radius of the planned flight path are irrelevant and alerts are received only when variables deviate enough from prior values to warrant considering a change in route. A VIRT data server accepts these COIs on behalf of the pilot and begins diligently monitoring for COI events. The pilot is not likely to see more than 5 alerts on the whole flight, well within his processing capacity. If each alert is accompanied by 100 data values (to update the display), the 5 expected alerts present about 100,000 times less data than communicated in the pull or simple push environment. These differences are significant and are very attractive to the pilot. While we don’t know exactly how this reduction in workload translates into mission outcomes, we can be sure that the pilot won’t be glutted, will notice the events, and will have sufficient attention resources to deal with them. Further, we have reduced bit flows by 99.999%, freeing up critical communication resources for other purposes.

When you improve bit flows by five orders of magnitude, you will change the organization in qualitative ways that will need to be supported by an appropriate system

architecture. Throughout biological and man-made systems, each order of magnitude (10 X) change tends to induce both structural and qualitative changes. The benefits of VIRT may not always be as great as five orders of magnitude, but they will be huge. In a recent thesis, using a discrete event simulation, my student LCDR Ray Acevedo [11] estimated a 200-fold improvement in bandwidth utilization when VIRT is employed within a Navy air tracking system called Cooperative Engagement Capability (CEC). The two orders of magnitude improvement means that considerably less communications bandwidth is required to convey important bits and every receiver has 99% more time to spend processing important information.

- **VIRT is a well managed information value chain**
 - Commander’s intent drives users’ COIs
 - Users’ expected COIs drive the course of action
 - Information suppliers analyze COIs and vie to serve them
 - User is served significant deltas from expectation, with high value per bit
 - User refines COIs and expectations
- **UDOP is a visual display of data sources available in common to the community**
 - Information available is not pre-determined
 - The user selects the data sources, value-added services, applications and presentation
- **A COP is a visual representation of a common database shared by some community**
 - The information available is limited to pre-arranged data sources

Figure 2. Comparing VIRT with UDOP and COP [9]. In a VIRT system, users define and refine expected conditions of interest (COI), while providers search for critical deltas and serve the users high value per bit.

These advantages are too big to ignore, yet there are many reasons to believe current DoD and DHS approaches won’t attain them without some change in direction. In the next section, we consider research and development needs and opportunities that can help these organizations deliver the benefits of VIRT to HFNs.

RESEARCH IMPERATIVES & AGENDA

The US DoD is implementing the Global Information Grid as a way of providing integrated information and information processing services throughout the military. The architectural approach aims to leverage Internet and web service technology, using a Services Oriented Architecture (SOA). The foundation services are called Net-Centric Enterprise Services, or NCES (see www.disa.mil/main/prodsol/cs_nces.html). The basic objective is to provide the infrastructure and tools for information superiority, where each operator can find and access all relevant information.

Beyond SOA, this goal of finding and accessing all relevant information requires that suppliers and consumers of information share some vocabulary and semantics. In the semantic web community, the term used to describe the concept of clearly specified semantics is *ontology* (see <http://www.w3.org/2001/sw/>). When users agree on what terms they use and what they mean, they have formulated such an ontology. I will also refer to this informally as the users' *vocabulary*.

So, users will refer to information categories and express COIs using the terms of an ontology. An ontology for pilot weather might include terms such as *icing*, *thunderstorms*, *dust storms*, *icing*, *visibility*, and *ceiling*. An ontology for pilot avoidance of air defenses might include terms such as *probability of detection*, *safe distance to avoid detection*, and *SAMs*. Users will be able to find relevant information among many possible sources, because each source will annotate its data with meta-data describing the categories of information included. In short, users can find relevant information by seeking sources that advertise available data with terms matching those in the users' ontology.

Such information advertising might be done in various ways. The DoD data-sharing directive [2] takes an initial step by requiring all suppliers of information to annotate data with XML tags associated with an XML schema that defines the vocabulary or types. Beyond this, many groups throughout DoD and the broader Federal Government have formed communities of interest to work out common vocabularies and XML schemas (cf. <http://it.ojp.gov/jxdm/>, <http://web-services.gov/>)/

In the business world, where similar efforts have been organized, such as for library resources (the Dublin Core, see <http://dublincore.org/>) and mortgage transactions (MISMO, see <http://xml.coverpages.org/mbaaXML.html>), efforts to build useful solutions take a decade or more. Practical ontologies result only when operators (end users of data) have their needs addressed. This requires considering the transactions or processes that will operate upon the information, because the value of information is determined by its ability to be used quickly, easily, and beneficially. Because information ontologies are not easy to create, failing to consider carefully how processes employ them to accomplish important transactions will squander time and money.

Our concern with current approaches to semantic mark-up and the development of information schemas arises from the failure to focus clearly on information integration in the context of processes that deliver VIRT to operators. Just because information is accessible and findable does not mean it's valuable. In fact, most mark-ups are useless

for finding and delivering valuable and timely information. While a community of suppliers, such as weather forecasters or logisticians, might develop a schema on its own, there's little reason to believe their semantics will prove suitable for operators or VIRT agents.

There are several reasons why supplier-oriented communities of interest won't likely give us the semantics and ontologies operators need.

- Suppliers don't know what problems their users are trying to solve or how they solve them.
- Suppliers don't use the same vocabulary or actual concepts as operators do.
- Operators' objectives, missions, and processes are evolving faster than supplier community efforts to standardize ontologies.

So, being brutally frank, we could easily spend a decade or more on efforts to standardize information ontologies **before** even beginning to evaluate the putative benefits of modern network and service infrastructures to deliver VIRT. At that time, we'd discover that operator COIs are the foundation for specifying value and that COIs and vocabularies vary among different missions, echelons, and roles. So at that point we'd find that we didn't have the required ontologies, couldn't specify operator COIs, and couldn't implement VIRT. In short, after a long infrastructure development, we'd realize we hadn't even begun to deliver valued information.

Rather than predictably proceeding into such a failure, we ought to consider if there's a better, lower-risk approach. I think there is a substantially better path, and that's what I want to describe now. Some of the steps on the path could benefit greatly from improved methods and technology, so I nominate those for a place on our R&D agenda.

The basic program I'd propose for delivering the benefits of VIRT follows:

1. Identify particular mission types to drive development.
2. For each chosen mission, define the process for effectively adapting to changes in assumptions, including the definition and monitoring of COIs.
3. Develop a vocabulary, semantics and related expressions for operators in those missions to define such COIs.
4. Create a data model suitable for representing dynamic values of variables needed to compute the COI expressions.
5. Determine sources adequate for populating the data model and computing the COI expressions.

6. Implement data collection, COI monitoring, and operator alerting.
7. Conduct after-action reviews to determine when:
 - a. COIs were evaluated incorrectly, or
 - b. a desirable alert should have been provided, but no appropriate COI had yet been defined.
8. Enable operators to define new or modified COIs in response to (7) and quickly evaluate these.

Each of these proposed activities can be implemented using existing technology. However, each of these steps could also benefit significantly from improvements in the underlying technology. The biggest opportunities for significant improvements that I perceive lie in these areas:

1. Provide a reusable modeling framework for mission types with explicit representations of their goals, activity models, assumed and predicted states, assumptions, and justifications. Begin to populate this framework with reusable elements from missions with common elements.
2. Develop a basic, tailorable process for monitoring COIs and alerting operators. Employ this process in various missions and continually improve it. Tie it into off-the-shelf technology components. Use it to help influence relevant standards efforts.
3. Create reusable vocabularies that operators find natural and useful in characterizing their COIs. Define an expression language operators can write and read to define COIs that uses their own vocabulary simply.
4. Provide off-the-shelf “cartridges” or “blades” for the most popular database products that make it easy to define models suitable for typical vocabularies, expressions, and COIs. In particular, provide standard solutions for expressions involving space-time intersections. Make it easy to “mix in” space and time dimensions to virtually any ontology.
5. Working backwards from the needs of evaluating COIs, develop meta-data specifications that can characterize required capabilities of information sources. Work with suppliers to enable them to publish data with matching meta-data schemas.
6. Conduct pilot and field experiments, iteratively improving the demonstrated and transferable capabilities.
7. Provide tools to audit information flows and to determine specifically “why” particular alerts occurred or “why not” when they didn’t. These tools help operators understand and debug their own information value chains.

8. Provide tools to modify the information value chains by fixing bugs in the vocabularies, expressions or COIs. Provide tools that help evaluate proposed changes by assessing whether and how well they would have improved past performance (using audit trails and regression test cases).

One key area of interest in this list has to do with the computation of spatiotemporal expressions for COIs. A typical COI of this sort might be: “the aircraft will arrive at the rendezvous point no later than 1700” or “the re-supply vessel will intersect the convoy at location X at time T.” These are simple examples of expressions that compute whether some dynamic variable, moving through time and space, will fall within some acceptable interval, itself perhaps derived from another dynamic entity’s changing state. Such time-space intersections occur frequently in military and relief operations, so they represent an important capability for COI computation.

Here are some general cases of events and expressions where we should aim to significantly improve off-the-shelf technology:

- a. Does an entity’s route intersect another’s range of capability (e.g., detection, weapons) at some time t ?
- b. Where is an entity expected to be and what area is included in its range of capability at some future time t ?
- c. What other positions and areas are possible, even if not currently expected?
- d. Will one entity detect another? With what probability?
- e. How long will an entity’s plan (e.g., planned route to destination) take? Will the entity exhaust any of its resources (e.g., fuel) before completing?
- f. Does the probability exceed P (e.g., 5%), that the ranges of capabilities (e.g., detection, weapons) of two entities will intersect (over the time remaining)?
- g. How probable is it that two entities will interact (e.g., collide, detect one another)?
- h. If two entities need to interact continually (e.g., remain in communication), will they?

Today’s dominant database systems provide poor capabilities for the efficient computation of time-space interactions. In addition, they provide no practical support for projecting future values from dynamic models, as occurs in dead-reckoning or forecasting. Probability computations, especially over time-space intervals, also would be nearly impossible to compute using conventional database queries. However, a combination of these capabilities—dynamic projection, intersections and

probable interactions—seems to play a central role in anticipating expectable problems and intervening to prevent them before they materialize. For these reasons, we should put improving these capabilities high on the VIRT R&D agenda.

CONCLUSION: CHANGE OF DIRECTION REQUIRED

We can probably all agree on a few important points. Timely valuable information can improve decisions and outcomes. Information is potentially valuable if it could improve outcomes, but to realize that potential the intended beneficiary has to receive it, attend to it, consider it and act upon it in time. In situations where human processing capacity is limited, where data glut is possible, and where communication bandwidth is limited, we must give priority to high-value bits. The only way to do this is to know which bits would materially affect the receiver. This requires understanding the operator's current state—the goals, the assumptions, the beliefs—so that contradictory events can be detected and quickly conveyed. In short, VIRT directly addresses the needs.

VIRT requires a focus on operators in the context of missions, because plan assumptions readily reveal what conditions of interest must be monitored. These COIs, in turn, reveal what semantic categories are required, and this reveals the operator's ontology.

VIRT suggests and makes possible an incremental and evolutionary approach to information superiority. We can implement VIRT for a small number of mission types at a time, covering the most important COIs first. This leads us to develop the vocabularies and ontologies incrementally. In each mission, we can deliver increasing value in proportion to the number of COIs monitored. Over time, we can extend the ontologies, COIs and missions. In short, we can implement incrementally and reap benefits incrementally. Moreover, we can recognize and address the essential evolutionary requirements from the outset, creating a process and technology base that can continually improve its coverage of missions, ontologies and COIs. In that way, we will achieve orders of magnitude better results with substantially lower costs, while realizing the goals of information superiority a whole lot faster.

By recognizing that model-based communication networks represent the ideal to which hastily formed networks should aspire, we gain a deep and practical understanding of essential requirements. Recall the questions posed at the outset. **Who** needs **What** information, and **How** does that information **Find** them? **How do we assure receivers are not glutted** by a deluge of low-value data and consumed

by attendant low-value tasks? The answers should now be clear.

The answers to the first question follow: Any operator whose mission can now be expected to fail wants the new information that invalidates the current plan's assumptions. That operator urgently needs to know that information. More generally, each participant most values information that he or she would respond to by changing course.

The answers to the second question follow: We should assure that operators receive the information they would most value while holding back redundant, irrelevant, and immaterial data. This assures they have plenty of time to attend to high-value events. It prevents the system as a whole from squandering scarce bandwidth and attention resources on low-value data processing tasks. As a side-effect, this approach increases the amount of those precious resources available for dealing with other important tasks, including investigating and planning for unprecedented events.

MCNs and VIRT can be implemented today, with off-the-shelf technology. However, much of the technology available is primitive and low-power relative to the opportunities that MCNs and VIRT create. Therefore, we suggest an R&D agenda that combines incremental delivery of VIRT applications, reaping significant ROI incrementally, with concurrent improvements in the underlying technology and methodology base. Those investments should produce huge improvements across a wide range of networks, thereby accelerating and amplifying the overall rewards. VIRT promises enormous increases in HFN performance as a result of communicating less.

REFERENCES

- [1] D. S. Alberts, J. J. Garstka, and F. P. Stein, *Network Centric Warfare: Developing and Leveraging Information Superiority*, 2nd ed. Washington, D.C.: Office of the Secretary of Defense (ASD/C3I/CCRP), 2002.
- [2] P. Wolfowitz, "Data Sharing in a Net-Centric Department of Defense," ASD(NII)/DoD CIO, 2004.
- [3] F. Hayes-Roth, *Hyper-Beings: How Intelligent Organizations Attain Supremacy through Information Superiority*. Bangor, ME: Booklocker.com, 2006.
- [4] F. Hayes-Roth, "Model-based Communication Networks and VIRT: Orders of Magnitude Better for

Information Superiority," presented at MILCOM-2006, Washington, DC, 2006.

- [5] P. J. Denning, "Infoglut," *Communications of the ACM*, 2006.
- [6] J. C. Hoag and F. Hayes-Roth, "Semantic Reasoning for Adaptive Management of Telecommunications Networks," presented at IEEE International Conference on Systems, Man and Cybernetics, Taipei, Taiwan, 2006.
- [7] F. Hayes-Roth, "Model-based Communication Networks and VIRT: Filtering Information by Value to Improve Collaborative Decision-Making," presented at 10th International Command and Control Research and Technology Symposium: The Future of C2, McLean, VA, 2005.
- [8] C. Oros, "Helicopter Information Awareness Module (I-AM): An Example of a Model-Based Communication Network (MCN) Architecture," presented at 10th International Command and Control Research and Technology Symposium, McLean, VA, 2005.
- [9] F. Hayes-Roth, "Two Theories of Process Design for Information Superiority: Smart Pull vs. Smart Push," presented at Command and Control Research and Technology Symposium, San Diego, CA, 2006.
- [10] M. Krieger, "Transforming the Way DoD Manages Data," presented at AFCEA, San Diego, CA, 2006.
- [11] R. A. Acevedo, "Valued Information at the Right Time (VIRT) and the Navy's Cooperative Engagement Capability -- A Win/Win Proposition," in *Information Sciences*. Monterey: Naval Postgraduate School, 2006, pp. 125.