Slides for the Presentation to AIAA Reno Conference: Technical Integration Technologies Panel Discussion Technical Narrative

Dr. William Henry Jones

October 20, 2003

The following pages provide a technical narrative that more fully elaborates upon the slide set prepared for the panel discussion on Technical Integration Technologies to be presented at the AIAA Reno Conference.

1 Slide pst_vgrf_0122 – Technical Integration Technologies: Current State and Proposed Direction

Technical Integration Technologies (TITs) are tools and technologies designed to facilitate and capture the technical exploration and development process. Despite the somewhat unfortunate acronym, the name is appropriate in order to identify the more comprehensive and complex nature of the technology and set it apart from the relatively easier efforts often known as Product Data Management (PDM), Product Lifecycle Management (PLM), Business Management Integration, E-Business, and the like.

The name Technical Integration Technologies also sets the area off from another area: that of Collaborative Technologies (CT?). Collaborative Technology uses the facilities of the computer and the network to allow technical (and other) people, groups, corporations, and agencies to work together effectively as a team. This team can be physically dispersed and may attach and detach elements as needs change. Information technology is used to keep the many details of such collaborations straight, to facilitate the transfer and review of information, and the like.

Technical Integration Technology, as I conceive it, goes beyond the goals of Collaborative Technology. Technical Integration Technology seeks not so much to integrate people into a functioning whole as to integrate the automated tools and information sources that those people traditionally tend and operate into a functioning whole. The machine is made to perform the day-to-day grunt work of turning the crank, keeping the details straight, following configuration changes, and the like. Because the tools themselves are made to cooperate, people are removed from that active role to a supervisory level where they are free to do what people do best: reflect upon the situation and apply what can only be called "judgement". The operational results of Technical Integration Technology will be the feedstock of Collaborative Technology.

This presentation will review the current state of Technical Integration Technologies and propose directions for the future of this area.

2 Slide pst_vgrf_0123 – Current State of the Effort

The Technical Integration Technology area is some 15 or more years old. As such, it is an area and technology of growth, as yet un-matured and un-settled. It is populated with both commercial vendors offering products for sale and numerous research efforts leading to or arriving at commercial offerings.

For the most part, the commercial offerings are man-in-the-loop efforts relying on a technologist to provide the organizing/integrating knowledge necessary to pull disparate applications and information sources into an integrated, cohesive, functional whole. Some offerings are centered about Computer Aided Design (CAD) systems since the object of much integrated analysis is physical in nature. Most offerings are, or are becoming, net-enabled, allowing many computers and technologists to cooperate and collaborate. There is a certain blur between collaborative and integrative technologies, and some of these products address both areas. In addition to being commercial successes in their own right, most of these products have scored successes in value returned to the customer.

At the same time as commercial products are being offered, research efforts exist and continue. Many different formulations and avenues of approach are offered. This state tells us that the issues of the technology are not fully explored, not entirely settled.

3 Slide pst_vgrf_0124 – Recurrent Stumbling Blocks

In my opinion, several stumbling blocks keep recurring in the various technical integration technology products and research efforts: a geo-centric focus or orientation of the technology, the retention of a human being in the integration loop, an unfortuante focus on the source of information rather than on the kind of that information, and an over-optimistic expectation that one source will solve the entire integration problem.

The geo-centric focus reflects an orientation toward a physical reality, something very natural in the manufacturing industry. In this selection, the geometry of the product becomes the center of all things: the shape of the fan blade, the thickness of the wing, so on and so forth. The physical design of the object/product is extracted, then analyzed, and the results associated back with that design. This is all well and good as long as there is a physical design, but what happens when there is no physicality? How does a pharmaceutical, a chemical, or a propellant company fit their "designs" into a CAD system in order to perform an integrated analysis? And if such companies are incompatible with aerospace companies because of their differing integration needs, how will the whole cooperate on an integrated design of, say, a Mars mission?

By man-in-the-loop integration, I mean making a person the agent by which the connections of integration are decided. Typical schemes require a person, operating through some computer-assistance tool, to connect the dots of an integration; this output from here goes to this input over there, this from here to that over there, and so on. This technology very quickly limits the magnitude of integration that can be achieved. Some put that practical limit at 15 to 20 applications. If such an integration requires 1000 individual connections to be made, a person working with the proverbial 99.99% accuracy (which seems somewhat optimisitic for people) has only a 90% chance of making all 1000 connections properly. Extend the task to 10,000 connections and the probability of an error-free result drops to 37%, and that requires you to believe that a person with 10,000 connections floating about in his head still works with 99.99% accuracy.

Another stumbling block of some integration technology is the focus on the source of information rather than its kind. This is reflected in the technologies that simply facilitate the acquisition of a number from a known

location in one output file and the placement of that number (sometimes modified in one way or another) in some other input file. This approach breeds an unfortunate inflexibility. Should another application come along that develops the desired number in a better manner – greater fidelity, lower cost, faster speed, and the like – the old connection must be broken and a new one formed. Furthermore, the introduction of each application or information source into an integrated suite necessitates that every other consumer of information with an interest in the offerings of the new element be "told" about the new-comer. When the integrated suite is small – 5, 10, or 20 elements – the effort can be managed, but as the numbers grow past 100, 1000, and beyond, the level of effort explodes to the impracticable.

A final difficulty with current approaches is that many providers expect that their offering will be the complete solution: their offering will wrap and integrate applications, information sources, design processes, and what not; visualize the results, produce design drawings, and machine control streams; inform every engineer everywhere when it is time to put in his or her two cents; spawn and feed every handy tool used by anybody; on and on and on. To the extent that any given solution does not achieve this vision, it limits the payoff that it can offer. Furthermore, this expectation leads to other difficulties.

1. Integration of the entire technological world is a huge problem requiring expertise in many disciplines beyond that of integration techology itself. It is unlikely that any organization will be so expert is so many fields.

2. The magnitude of the problem is beyond the likely capital resources of any single entity. The single-source expectation means that each provider must solve the entire problem. This diffuses capital investment rather than focusing it.

3. Beyond the size of the integration problem, it is unlikely that any single "look and feel" will meet every need; the multiplicity of products in any marketplace reflects the fact that different people and companies like to do things different ways. The single-source is antithetical to this reality.

4 Slide pst_vgrf_0125 – The Fatal Flaw

The single most glaring flaw in the current state of affairs is that, in a business/technology area whose sole purpose is to make things work together, the products and systems offered do not, themselves, work together!

The vaguely conscious customer will immediately understand these things.

1. The vendor does not understand the product he proposes to supply. In a business in which making things work together is the entire point, the vendor's product doesn't work with anything else.

2. No matter what product the customer picks, his company/organization will only be able to cooperate/integrate with others that picked the same product.

3. The customer will recognize this as the same disastrous path he has lived with his CAD vendor choice.

   Virtually every company and organization on earth has had the experience of being unable to share CAD drawings/information electronically with a partner. Some large organizations have experienced this difficulty within their own organizational units. Many have had to solve the problem by forcing the weaker partner to buy the stronger partner's choice.

4. The customer will also understand that the future fortunes of his company will be guided by a choice he must make now based only on his anticipation of future trends and results which he can hardly be expected to accurately foresee. Should he chose poorly, his company will incur the expense not merely of a back-track on vendor choice, but also of migrating/re-doing all of the work performed with the incorrect choice. Imagine the cost of re-doing all of a company's CAD drawings on a different vendor's system, and then multiply that by re-doing all of the analysis and validation that goes with the products described by those drawings.

5. The customer will understand that, even should he chose correctly, his company's capital assets of information and capability will become the hostage of his vendor. While he may expect that the rules of mutual avarice will lead to continued, productive business arrangements, he

still runs the risk that his particular needs may not be sufficiently addressed by those arrangements.

Very large customers can usually expect that their needs will drive their vendor's offering, but many smaller groups have seen their needs neglected because those needs did not have a sufficiently broad appeal to justify the vendor's capital improvement effort.

All of this combines to leave the customer in a terrible state of indecision and reluctance. Furthermore, the aggregation of negative factors works against the cost, time-to-market, and performance gains needed and desired to secure a commanding competitive position.

5 Slide pst_vgrf_0126 – The Reasonable Solution

The clear solution to the fatal flow is for all of the participants to adopt/move toward an open technical integration technology standard. This standard should be a foundation layer encapsulating the core technologies of information representation and application integration; that is a common foundation upon which all can build.

An open, industry standard would have the following significant benefits.

1. The adoption of an integration foundation standard would convert the situation from one of may different, isolated vendors running about in the marketplace to one of many cooperating and coordinating partners participating in a single industry. This would replace the cut-throat world of competitive advantage with the more stable and predictable world of co-operative advantage.

2. The vendors of technical integration technology products would send the clear message that they "get it", that they truly understand the nature of their product. Their various products would interoperate within the standard-based world, co-operatively solving the many aspects of the technical integration technology puzzle.

3. The customer's issue of which way to go, and whether to go at all, is eliminated. The customer would not have to worry about whether or not his future partners in larger enterprises would have gone another way. The "industry" would have spoken and the standard would be the one and only way for everyone to go. The door to all customers would be permanently opened; many could sell pieces at many doors where before only one could sell a whole at any given door.

4. Vendors within the industry would no longer have to solve the entire integration problem within the limits of their own capitalization. Instead, they could devote their entire resources to solving some particular part of the integration problem to such an extent as to become a commanding presense in that particular area. For example, one or two vendors could focus on visualization products while another grouping dealt with developer's workbench products and another with data mining and fusion technologies.

Further, working pieces to the solution to a standard extends the reach of the capital assets of the integration industry by reducing redundant solutions. Because each participant or competitive group solves only a portion of the puzzle very well, rather than the entire puzzle poorly yet again, the overall result is a far better solution to a very demanding problem.

Additionally, the presense of a standard may draw in the participation of second- and third-tier participants. Vendors of design and analysis tools might be more tempted to devote their own capital resources to the integration aspects of their offerings in the presense of a standard-based industry than in a marketplace where many solution strategies still exist and compete. Also, a standard-based industry is likely to foster consultants and other small support firms with only sufficient resources to address small niche opportunities. While each of these contributions might be discounted individually as too small to be of significance, the aggregation might well significantly extend the captial resources of the integration industry and positively influence the overall result.

5. The customer can participate in the solution of the integration problem as his focused needs dictate. Should the customer's needs be so narrow that no vendor is willing to devote resources to their solution, the presense of an open standard provides the customer a reasonable expectation of being able to provide his own solution or of finding a niche firm will to take on the task for him.

The simplest expression of the bottom line is that it is better to be Intel, Advanced Micro Devices, AST Research, Dell, Compac, Gateway, HP, IBM, Micron, Microsoft, Borland, Intuit, *et al, ad infinitum*, than to be Apple Computer or, worse yet, Digital Equipment Corporation.

6 Slide pst_vgrf_0127 – The Nature of the Standard

In general terms, the nature of the standard is easy to define.

1. The single most vital issue is to secure the standard as being a standard; the slightest hint of incompatiblity between "versions" will have the most dire effect. This is most likely to be achieved by a documented, open-source, freeware rendering of the standard as actual compilable software, probably administered by an industry consortium.

   The traditional standards approach of a document to which many vendors can build software implementations is far too likely to leave cracks through which the goal of interoperability will slip through. It must always be remembered that the entire point of the integration industry is interoperability; each possibility for falling short of that mission diminishes the industry.

2. The standard must be net-enabled. In this day and age, anything else is unthinkable.

3. The standard must be operationally secure. The expected focus of integration technology is upon valuable business resources and processes; information that can be put to profitable use, applications that analyze problems and provide value-added answers, technology explorations that must leave an auditable trail for certification and approval, and the like. The security of such activities must be at least believable or customers will certainly go elsewhere.

4. The standard must scale to the limits of imagination. Not only will this technology enable the practical implementation of such things as bio-informatics, which measures its data volumes in petabytes, but it will indeed be a cause of such data volumes as integrations of complete, complex systems analyses such as entire air vehicles and inter-planetary missions become possible.

5. The standard must exploit the technology of self-revelation. No single, rigid formulation of applications and information could possibly embrace the entire diversity of technological investigation and analysis.

6. The standard must exploit the technology of semantic infusion to allow consumers of information and application to focus on the kinds of information and services desired rather than the particular sources where such information might be obtained. This eliminates the N^2 integration problem since the kind of information can be recognized regardless of its location.

7 Slide pst_vgrf_0128 – The Direction Beyond the Standard

Given a standard for self-revealing applications, information representation, and integration and given a sufficiently rich environment (that is, an environment that actuallly has the resources to solve the problem to be posed), it should be possible to devise an algorithm which, itself, can devise a method of solution for problems posed as the desire for the best achievable value of some parameter or combination of parameters. Then, given that devised method of solution, that method can be exercised by something amounting to an optimization process until the best achievable result is identified.

In simple terms, this algorithm for autonomous solution is like a program linkage editor: a search is made for an application producing the desired result. having found such an application, it is asked what inputs it needs to produce that result. A recursive search is conducted to find applications producing those pieces of information, each recusion then producing the need to find more outputs to produce the currently needed inputs. The algorithm completes when all the required inputs are of such a nature that they may be treated as the elements of an independent design vector. The reality of this algorithm will be, or course, a good deal more difficult.

The benefit of this technology it that it takes that final step beyond the manual integration of 5, 10, or 20 applications to enable the intergration of hundreds and thousands of applications into a single, functioning whole. The man may be befuddled by ten or twenty thousand parameters connecting here, there, and everywhere; for the machine, it is just another day at the office.

8 Slide pst_vgrf_0129 – And, As It Just So Happens...

The Project Integration Architecture (PIA) is an application/information integration framework that meets all of the requirements and goals previously touched upon. It develops and exploits the technologies of self-revelation and semantic infusion through class derivation to allow consumers of information, both man and machine, to deal with any such resource on a discovery-oriented basis.

PIA is presently implemented using CORBA-served, distributed object technology. Per-object, per-user access controls assure security while a distributed lock mechanism resolves concurent access issues for a full, multiple-accessor environment. Objects may be served either by single or clustered servers. Clustered servers will soon support automatic fail-over in the event a member of a cluster is shut down in an orderly manner.

PIA servers and server clusters can be connected into a larger grouping called a collective. Any entity logged in to any member of the collective enjoys access to all of the information of the collective, subject of course to access controls placed on individual information resources. Given the great diversity of information and services that may be accessed through PIA, this opens the door to a single log-on to enterprise-wide, and even mutli-enterprise resources.

The persistence mechanism supports redundant storage options for improved reliability. User-directed storage location is also supported for those situations in which a user wishes to provide his own storage services rather than use those of an out-of-house provider.

The object "address space" of PIA extends to a practical infinity. Presently, every machine on the internet could be devoted to producing 16 billion billion objects every second for the next 500 billion years before that space would be exhausted. Objects may be of any practicable size; however, the extremes, both large and small, are less advantageous.

The technologies of self-revelation and semantic infusion through class derivation exist and have been demonstrated to function as proposed in information propagation activities. Applications have acquired information from other applications based upon the kind of information desired, rather than the specific source of that information.

In a sense, semantically-infused parameter objects become self-aware. For

example, parameters that are dimensional in nature (that is, parameters that are defined in some system of measurement) are aware of their dimensionality. This awareness extends both to their dimensional characteristics (for example, a velocity being in terms of length per unit time) and the system of measurement in which they exist (for example, feet per second). Such dimensional parameters will not allow accesses that are not, themselves, dimensionally aware; that is, one cannot find out that a measurement is simply 15, but must inquire so as to find out it is 15 feet per second, or 4.572 meters per second, or 5 yards per second, or whatever.

The infusion of dimensional awareness is, of course, but an example of semantic infusion. Parameters are able to infuse all manner of functionality appropriate to their kind. As another example, geometric solid parameters are able to include in themselves the ability to generate arbitrary cross sections. Flow field parameters would be able to provide derived quantities that are beyond the level of the raw data.

PIA is believed to be relatively platform insensitive; however, due to limited resources, it has not recently been demonstrated to be so.

The great body of PIA software is limited to the application/information foundation framework. A skeletal, Windows-based Graphical User Interface (GUI) exists solely for test and demonstration purposes. This GUI is kept deliberately minimal so as to leave the greatest possible opportunity for private interests to develop and profit from the great variety of front-end technologies that might be applied to PIA. Other products areas such as developer workbenches, data fusion and mining technologies, autonomous solution system algorithms, optimization packages, and the like are deliberately untreated, again, to leave the greatest opportunity to private interests.

PIA provides facilities to allow the entire path of technological exploration to be tracked as a structure of project configurations. This is desirable in fields such as spaceflight in which such a history may be a critical record after a failure. In some industries, notably the pharmaceutical industry, such records are necessary in order to obtain certifications and approvals necessary to market a product.

When two or more application co-operate in an analysis of a technological process, their project configurations are kept in rigid synchronization. This simple matter of bookkeeping assures that un-manufacturable designs (for example, a boiler design with the weight of the 16-guauge shell and the strength of the 1-inch boiler plate shell) never reach the shop floor.

PIA further recognizes that different configurations of a technological pro-

cess can represent independent problems that may be pursued concurrently and, as an option selected by default, PIA will spawn such parallel operations off as independent "threads" of execution. While a human may become confused when trying to keep the concurrent analyses of five or ten different variants of a design straight and sychronized with the work of others, PIA's machine-implemented bookkeeping has no problem keeping hundreds or thousands of such concurrent analyses, spread throughout the computing resources of serving clusters and the greater collective, straight and documented.

Beyond the project configuration capability, PIA provides a diverse set of other facilities to track operations and efforts in other manners. Currently, a digital signing facility is contemplated that would allow the state of an aggregation of applications to be established at arbitrary times. This is expected to be a key tool in resolving intellectual property rights issues when such issues arise.

In the recent past and at present, the Government has made PIA available in source code format at no charge to those willing to sign a Software Use Agreement. Two companies are presently pursuing product developments based upon PIA technology. The Government is still considering the issues involved in making PIA available as freely-downloadable software.