

Open Architecture Testers – Year One: Does Reality Match The Vision?

By John Johnson, Don Edenfeld, Jim Neeb

Intel Corporation

Intel's test equipment roadmap is complex and very difficult to manage. We have over 1,000 high-end automatic test equipment (ATE) testers running more than 1,000 products across a worldwide factory and engineering site network. The costs associated with purchasing, support, training, etc., run into multi-billion dollars and represent a significant portion of product test costs. In addition, most of these platforms are five to 10 years old, making the fleet increasingly costly to maintain.

Our products demand increasing test capability over time. In many cases, we would identify a minor change to improve capability that could significantly extend the useful life of a platform. However, equipment suppliers are rarely motivated to extend the life of older products and such improvements require the manufacturer's expertise to implement. Intel found itself with a tester fleet that was becoming obsolete for frustratingly simple issues. Changing test requirements resulted in extensive efforts to select, develop, integrate and ramp a new platform.

New platforms have been added to the fleet on approximately a three-year cycle over the last 10-plus years. With a variety of platforms from different suppliers, we expend extra effort in managing multiple platform operating systems and processes for test program development. Some products must enter manufacturing on multiple platforms to enable test coverage or sufficient test capacity. This environment is increasingly resource-intensive for test content development and management, but it is justified as a lower-cost alternative to fleet replacement strategies.

We looked at our upcoming ATE challenges: new products with orthogonal test capability requirements, the need for increased capacity across new and existing product lines, and a very strong requirement to reduce manufacturing test costs for every product. It became clear to us that simply following our traditional selection and procurement processes to identify another traditional test platform was not sustainable.

The Vision

Intel's test community explored the problem and formed a list of the ideal attributes of a new platform. Any new platform would need to:

- Have a low base cost to support low-cost, low-margin products;
- Keep the high-value instruments and their associated IP in a modular form to allow very flexible and targeted configurations;
- Support a technical base that could scale to support high-end, high-performance instruments;
- Be fungible with a long sustainable life and be able to host new instruments that had not been conceived;
- Be applied broadly and support a common tester OS to leverage a sustainable development environment;
- Be built on a standard architecture that allowed us to build our own solutions incrementally as we required them; and
- Enable multiple supply sources for best-in-class purchase selections and competitive pricing.

This list became our vision, and in 2001 we set out to survey what was available

in the industry. What we found was a very strong push by virtually every major ATE vendor to standardize their platform offerings on their own proprietary version of a common infrastructure architecture. These platforms would employ modular instruments that plugged into "universal" slots. Their motivation was to consolidate their engineering efforts and reduce costs by focusing on the high-value portion of the platform: the instruments. They would leverage the infrastructure across multiple products and compartmentalize their engineering efforts.

Their solutions addressed some of our vision but missed on key points:

- The lifespan of these platforms is dependent on the supplier's product roadmap, financial health and business decisions.
- Most of these solutions still left IP in the infrastructure (instead of pushing it all into the instrument), which limited future applications and contained expensive overhead with no obvious end-customer benefit.
- A customer would remain dependent on a single supplier's product capabilities, schedule and pricing.

Driving The Open Architecture Vision

At Intel we realized our goals would not be met with a proprietary test platform, and the path seemed obvious – only an open modular architecture would fully meet our requirements. Standard architectures had already proven successful in bench and rack instruments incorporating standards such as VXI and PXI.[1] Intel was motivated to adopt an open architecture platform because it addressed our challenges and vision

directly. A platform designed to host complex ATE instruments via a standard interface could incorporate new capabilities incrementally as they are developed. With multiple suppliers participating, competitive pricing and features would emerge. Additionally, the configuration of these platforms could be cost-effectively tailored to specific test requirements and changed as those requirements evolve.

Intel and a handful of other companies joined forces in 2002 to foster an ATE open architecture and formed the Semiconductor Test Consortium[2] to promote open specifications. These specifications define the standard hardware and software interfaces that are used to implement the tester architecture. Platform development ensued, and Intel began proliferation of open architecture test platforms into both manufacturing and engineering sites in the fourth quarter of 2003. Internally, we call our growing fleet of open architecture-compliant systems “CMT,” for configurable modular tester.

Since that time we’ve gained experience with both the benefits and challenges of CMT platforms in a variety of situations. After one year, a review of the goals provides some insight into our open architecture experience.

Low base cost: Intel has placed CMT into applications that are also supported by our dedicated low-cost structural tester. On the basis of cost and capability, a properly populated CMT competes directly with the dedicated, optimized and low-performance structural test platform.

Modular: The CMT has a common slot architecture, a very flexible instrument communication bus and an interchangeable connection layer to the device under test. In addition, the tester operating system software is a modular, extensible architecture and is designed to allow new capability additions.

High performance: CMT instruments in use or under development support the highest frequency/bandwidth RF products, the fastest logic serial bus interfaces and the highest power microprocessor products. There are no foreseen technical limits for new

solutions. However, the open architecture specification is not static and will evolve over time as the committees identify improvement areas.

Long life: One year is not much to go on; however, no issues have surfaced to give us concern.

Applied broadly: We have encountered significant “pull” from our factories and end users. In many cases, we have had to slow adoption of CMT platforms to ensure our development and support resources were not stretched too thin. Virtually every product line at Intel is aligned to current or future use of the CMT platform.

Standard architecture: The open architecture specifications are published and available to consortium members. Hardware and software developer kits are in place. The review and ratification process is functioning.

Multiple sources: A handful of suppliers are developing and releasing open architecture solutions less than one year after the specification set was published. Intel begins placing third-party supplier solutions in our factories early next year.

fundamental requirement. Some of the restrictions previously imposed by the supplier must be re-imposed by the user, but can be more intelligently optimized for a particular customer’s environment and product mix.

One common case where configuration management is important is in situations where device locality is critical. High-performance signal and power delivery require a very short and highly optimized path. Special care must be taken to define instrument population schemes that enable large varieties of current products to coexist while maximizing headroom for future requirements.

As we encountered these variations, it became clear that we needed an internal consultation and configuration management review process to ensure different configurations are as fungible as possible. Proposals are now reviewed with respect to current and future product needs as well as existing tester configurations.

Commercial agreements are affected in ways we had not anticipated. When we procure traditional platforms, they are treated as single entities from a specific supplier. We use routine processes for measuring supplier commitments to system performance, capability and price. However, this

platform-level assessment is no longer meaningful for a modular platform with instrumentation sourced from multiple suppliers. Other approaches are required to manage the open architecture environment.

A simple example of this complexity is in measurement of mean time between failure to guarantee equipment availability. Rather than simple system-level measurement and tracking, we must measure equipment performance by breaking down the platform into its modular components. Because of the large number of individual instruments (thousands), the task is large. Even simple concepts such as instrument warranty period get confused. We are defining new methods and commercial terms to deal with these challenges.

Software in a modular open architecture environment is a significant challenge and >>

A handful of suppliers are developing and releasing open architecture solutions less than one year after the specification set was published.

From this experience, we conclude that the vision is quite solid. The proliferation of CMT into multiple applications continues to accelerate. While our critical problems appeared to be resolved with the modular open architecture platforms, we have also seen new challenges emerge. Considering the very short development cycle – concept to equipment ramp in 1.5 years – there are the obvious startup issues to be resolved that come with any new platform, and we are not completely finished with that effort. It is the new attributes of the platform, however – dealing with flexibility, modularity and multiple suppliers – that have presented the most interesting challenges.

Modularity is a great benefit, but it also introduces significant complexity. To maximize manufacturing efficiency, configuration compatibility is a

>> configuration management requires careful attention. Consider that any tester operating system consists of hundreds of thousands of lines of code spread across more than 1,000 files. With open architecture, that core OS then needs to interface with multiple versions of instrument drivers and modular language and tool extensions. The many resulting combinations of software and hardware revisions across a large global network of factories cannot be managed manually.

Factory test programs and developer environments need to remain very stable to prevent excursions and surprises that result in increased cost and lost productivity. It is important that a baseline “qualified” state can be

group is consistent and flexible parallel device test methods. Accounting for all the combinations of both distributed and shared instrument resources is a fundamental challenge that has never really been resolved in a systematic way on any previous platform. Quick fixes are not acceptable, and true open architecture standards demand solutions that anchor diverse applications over time. As a result, we anticipate a few more cycles of software rework as the working groups blaze new trails.

Conclusion

This article presents a few of the lessons that have emerged since we first captured the original goals list. Three

One focus of the STC's SW working group is consistent and flexible parallel device test methods. Accounting for all the combinations of both distributed and shared instrument resources is a fundamental challenge that has never really been resolved in a systematic way on any previous platform.

preserved. However, the same tester that is used to run a stable product is also used to qualify untested software features and new modular components. To solve this problem, an installation and configuration management system was developed to encompass both hardware and software state choices. Compatibility checking capabilities allow sharing multiple configurations that are acceptable to a particular product test program but eliminate unwanted or unknown differences.

Obviously, the software base is evolving with refinements in open architecture capability. A recent “refactoring” exercise was done to improve interface consistency between the module control software and the core OS. The open architecture extensibility model is fundamentally sound and has been applied successfully to new digital test feature additions. Work is under way to establish similar framework for analog and RF instrumentation.

One focus of the STC's SW working

years is not a particularly long period for such an ambitious undertaking, and we anticipate more challenges ahead. Intel has made significant progress with our CMT platform proliferation, and we have a growing list of applications for our fleet. There is much to do in continuing to refine the open architecture specifications and enabling tools. The list of instruments that need to be developed is extensive, and the supply base is still immature. Everyone involved in the consortium effort is on a steep learning curve.

Does our experience thus far show that the vision for a tester open architecture still viable? Inside of Intel, the answer is a resounding YES! ■

Endnotes

1. VXI = VME eXtensions for Instrumentation, VXIbus Consortium information at www.vxibus.org/; PXI = PCI eXtensions for Instrumentation, PXI Systems Alliance information at www.pxisa.org/.
2. Advantest, Motorola (now Freescale) and Intel founded the Semiconductor Test Consortium in April 2003. Information at www.semitest.org/site.

Biographies

John Johnson

John Johnson is a principal engineer in Intel's Test Capital Equipment Development group in Chandler, Ariz.

His career in ATE started 29 years ago as a maintenance technician on Air Force avionics test sets and then transitioned to LSI testers as a civilian.

For 15 years, he was a test engineer with responsibilities in device design debug, manufacturing test, fault isolation and product health. For the last nine years he has led Intel's programs to define and develop new manufacturing tester architectures.

Jim Neeb

Jim Neeb is a staff test engineer in Intel's Test Capital Equipment development group in Chandler, Ariz.

After graduating from South Dakota State University with a BSEE, he started his career 14 years ago as a product engineer working on Intel's CPU and chipset products.

For the last seven years he has defined and developed numerous manufacturing test equipment solutions. Jim holds five patents in test methodology and tester electrical/mechanical/thermal hardware design.

Don Edenfeld

Don Edenfeld is an engineering manager in Intel's Test Capital Equipment Development group in Hillsboro, Ore.

He received a master's degree in electrical engineering from the University of Virginia in 1996. He is the chairman of the Test International Technology Working Group for the International Technology Roadmap for Semiconductors.

His team is responsible for equipment strategies for engineering and production test.