AUTOSAR AND FLEXRAY: A TALE OF TWO STANDARDS

SERGE LEEF, MENTOR GRAPHICS

Mentor Graphics®

A U T O M O T I V E

W H I T E P A P E R

The emerging automotive design software standard known as AUTOSAR (Automotive Open System Architecture) began as the product of an industry-wide effort among European auto makers and their suppliers. Its objectives are similar to those of software standards in other industries: to bring structure, clean interfaces and implicit methodologies to a process—in this case, the design of distributed systems within automobiles.

FlexRay™ is a serial bus communication standard that has evolved over roughly the same time span as AUTOSAR. FlexRay came into existence as a solution for the shortcomings of the prevailing automotive bus standards, particularly the CAN protocol. Like AUTOSAR, FlexRay counts many prominent automotive OEMs and suppliers among its advocates. Boasting much higher performance (in every respect) than other in-vehicle buses, FlexRay alone is suited for “x-by-wire” applications that must deliver absolutely predictable results for steering, braking, and so forth.

What do these lofty aspirations mean to the designer who needs to get a complex array of automotive functions working together with high reliability? To the executive responsible for minimizing costs while delivering timely, compelling products to customers? To the end-users of tomorrow’s automobiles?

Clearly the automobile industry is in transition. Government regulators, if not auto buyers, are pushing for smaller, cleaner cars—an evolving imperative that equates to man-years of engineering time spent on research and development. Add to this the burden of electrical/electronic distribution, growing more complex with each new feature that either regulators or consumers demand. And of course there is unending pressure to improve reliability and expand design flexibility while reducing development costs.

Never was there a better time to adopt technologies that can simplify the designer’s job and streamline the end products that embody his or her work. AUTOSAR is such a technology.

AUTOSAR is essentially a set of standards encompassing interfaces and software module definitions. Most importantly, it sets forth a structure for the embedded software that ultimately operates a vehicle’s complex network-based distributed system. At the conceptual level, AUTOSAR can be viewed as a new platform that enables designers to focus on unique, innovative functions while insulating them from the implementation details of integration.

In the context of this discussion, the “functions” in a vehicle include climate control, transmission control, anti-lock braking, suspension, and many others both large and small. They encompass every subsystem that must interact with either the vehicle’s human operator/occupants or its other electrical/electronic subsystems. Today these features operate under the control of a phalanx of single-function Electronic Control Units (ECUs).

In the world of AUTOSAR, things are different. Functions consist of one or many software components (SWCs) that can be executed on any AUTOSAR-compliant ECU. The SWCs implement distinct, differentiated features in the end product.

Compare the AUTOSAR approach with an existing architecture made up of single-purpose ECUs. In the older approach, each ECU is associated with a fixed software definition. It is custom-programmed and hard-wired for one specific task. In contrast, an AUTOSAR-compliant ECU is an uncommitted “dock” for software components whose functions define the ECU’s role in the vehicle’s operation. Given two outwardly identical hardware modules, one may run the climate control system while its twin runs the dashboard functions.

AUTOSAR SWCs interact via an Application Programming Interface (API) known as the AUTOSAR Runtime Environment which fronts underlying mechanisms that carry out computation, control and communication activities. Figure 1 depicts a simplified overview of this model as it appears within an AUTOSAR-compliant ECU.

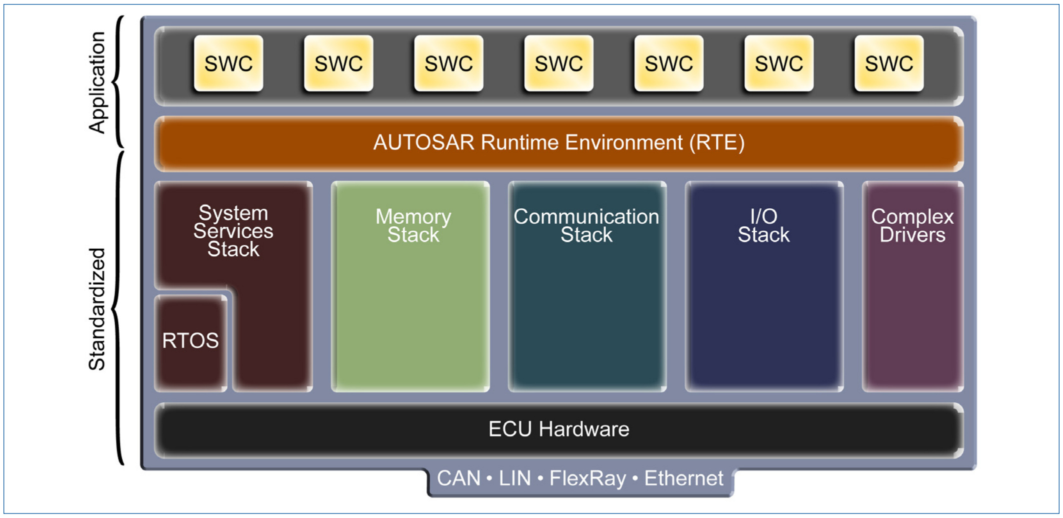

Figure 1: An AUTOSAR-compliant ECU loaded with standardized general-purpose software stacks that interact with SWCs to implement vehicle features.

The lower tier is essentially "generic," while the upper tier of SWCs is given over to the innovative spirit of the designer. What is the implication of this dichotomy? While today's typical car has 20 or more distinct and different ECUs, AUTOSAR may shrink this to just a few different ECU types whose behavior is determined by their application content. These types will occupy a hierarchy of performance that matches the jobs they will do; for example, a low-cost 4-bit ECU might suffice for door locks and window controls, while a 32-bit unit with large memory capacity would be needed to execute real-time power train management duties. While the car may still house 20+ ECUs, there might be just four or five different hardware types in this population.

While the ultimate AUTOSAR vision is in the initial adoption stages, it has the potential to significantly change the relationship between the subsystems in a car. And it will change the face of automobile architecture and design as well.

ENABLING A NEW APPROACH TO AUTOMOTIVE DESIGN AND PRODUCTION

If that latter statement seems like a grand claim, contrast the traditional way of automotive system design with the approach that AUTOSAR will enable. Most automotive brands are actually OEMs— known in other industries as "integrators"—that design vehicles and assemble them from components bought from tiers of contracted suppliers. One supplier might deliver the anti-lock brake system, another the climate controls. These functions are completely self-contained in that they include all the mechanical, electrical, electronic, and software ingredients required to implement a function.

It is a business model that makes it difficult and expensive for OEMs to change any aspect of a design after signing off an order with the contractor. Even a modest software change requires a new Request For Proposal, with all the attendant time and cost. With market economics pressuring car makers from one direction and exploding complexity pushing them from the other, much more flexibility is needed.

Therein lies a key benefit of the AUTOSAR standard: it allows OEMs to recapture a greater degree of control over the architecture of their vehicles. Assuming the existence of a diverse market in AUTOSAR-compliant components, an OEM can simply order the elements and standardized ECUs for, say, a suspension package from a specialty supplier. The OEM's design team then adds intellectual property in the form of the software components (applications) that make the suspension system deliver the comfort, safety and responsiveness that distinguish the car's brand.

OPENING THE DOOR TO OPTIMIZATION

At this point it's clear that AUTOSAR offers some direct efficiencies and economies that can be of great value to auto makers. But the flexibility of AUTOSAR's SWC concept promises other benefits that may pay even greater dividends in the long run. These are innate in AUTOSAR's open architecture:

1. There is no hard and fast rule that says one ECU must contain all of the SWCs associated with a particular function. Every ECU contains the same standard interface (namely, the AUTOSAR RTE) that can execute SWC code regardless of its role in the vehicle.
2. Many ECUs have some computing bandwidth to spare. Even in a world that demands efficient use of every internal computing resource, some ECUs will very likely have room for a few more SWCs.

Taken together, these two simple facts open the door to an optimization approach that can dramatically reduce the total number of ECUs required to operate an automobile. If a function such as cabin lighting requires five SWCs, for example, three of those might reside in the local ECU allocated for that job, while the remaining two reside in a nearby ECU that manages dashboard displays. Optimization can reduce complexity, save space and reduce cost, simplify wiring harness development, and improve reliability. In theory optimization will allow designers to configure vehicles with the fewest possible ECUs and the least complex wiring.

But will it work?

Optimization must be pragmatic, and pure simplification can't be its only goal. Consider this working definition of optimization:

> *"Optimization: manipulating a particular system's state such that it achieves a new, improved state that is better by objective metrics."*

These metrics reflect the goals of the design project, which typically encompass not only simplicity but also cost, performance, functional priorities, and many other issues. It might seem cost-effective to use a brake system ECU's spare capacity to support the air conditioner controls, but that risks compromising the braking response at a critical moment, and might even add cost to other components such as the wiring harness. In the real world, extensive analysis must be applied in conjunction with optimization.

ANALYSIS

To be a candidate for optimization, a system must have certain properties: an initial state that is easy to describe to the computer software that must process the information; variables that are easy to manipulate; and outcomes that are easy to observe and compare as they evolve.

Knowledge of the initial state of the system is key to the whole process. There can be no measurable ***after*** if there was no ***before***. Until now there has been no uniform way to characterize the state of an automotive system. Terms like "quality" and "efficiency" apply, but they can be rather intangible and subjective.

Traditional (pre-AUTOSAR) optimization has relied on the engineer's experience and intuition, conditioned by time, cost and legacy constraints. But the increasing number of variables (even if the consideration is limited to ECUs alone) geometrically multiplies the number of solutions that must be evaluated. The scale of the optimization space soon exceeds the grasp of human faculties.

AUTOSAR introduces discipline in the form of a language that allows designers to concretely represent a distributed system, manipulate it, and then assess the result. AUTOSAR's XML files are declarative files that comply with AUTOSAR formats and enable designers to describe systems, ECUs, SWCs, mapping, timing, and more; everything there is to know about the system in question. A climate control system, for example, might be fully defined by a dozen files relating to its overall functional behaviors.

In AUTOSAR-based design as in other creative efforts, the starting point or initial state is frequently a modified predecessor design that incorporates new features proposed by the architecture team. This "device" usually includes some intuitive SWC/ECU mappings which may or may not be usable as the design evolves.

Optimization is based on analysis, of course. There are two ways to approach analysis: static and dynamic. The former is more simplistic and more common at present.

In static analysis, the solution might begin with system representation that symbolizes a group of SWCs mapped to ECUs. The question to be answered is straightforward: Is this configuration good or bad? Is it appropriate to put all of the climate control functions, for example, into single ECU?

An AUTOSAR-based analysis tool can provide pragmatic answers because:

- It knows how much memory the chosen ECU has, and how much memory the climate functions require.
- It knows how many types of control messages have to be transmitted, received, interpreted, and executed at what rate; and it knows the throughput of the ECU's network controller.
- It can anticipate side-effects on other networks in the system.

It is important to note that static analysis still isn't an optimization procedure. But its findings can drive the optimization steps, and it allows the design team to develop and test basic configurations easily. A static analysis confirms that, yes, module "A" can receive and process packet "B" and furthermore that packet B is of an appropriate type.

The next step up in capability is called dynamic analysis. This approach incorporates simulation. Dynamic analysis begins with the construction of a set of software-based models embodying the best approximation the system elements' behavior. The object is to force the models to do just what a complex hardware/software combination would do.

Dynamic analysis brings the dimension of time into the evaluation. It might reveal, for example, that identical packets "C" and "D" arrive at the same time as packet B, and the module has no way to deal with these three events at once. Dynamic analysis is a more rigorous and realistic test of a design.

OPTIMIZATION TODAY…

Good analysis tools, static or dynamic, lay the foundation for optimization. Analysis establishes a baseline of performance values and/or other variables for a given design implementation. Manipulation is the next step, wherein the design is modified to correct problems and evaluate alternative concepts.

In a sense manipulation is half of a "cut and try" approach, though of course the integral analysis steps help the designer make informed choices about what to try from one iteration to the next. The manipulation is not automated; it depends on the user's theories, hunches, preconceived notions, and astute guesses. Figure 2 summarizes the iterative role of manipulation and analysis in the optimization process. It is important to emphasize again that AUTOSAR analysis tools can be an enabling element in this routine.

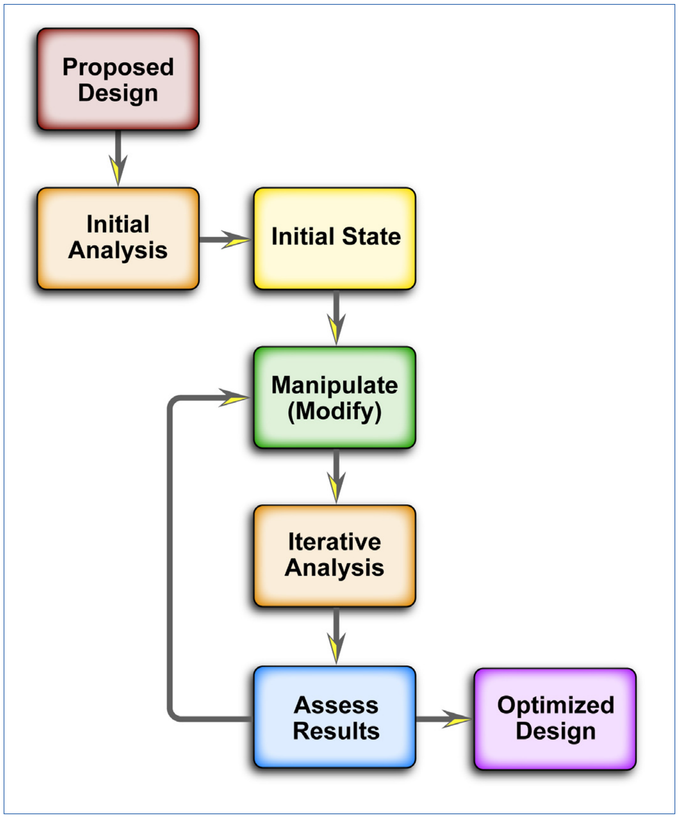

Figure 2: This flow depicts a manual optimization process. The automated approach embodies essentially the same steps.

To further illustrate the process, imagine an engineer who has just completed the initial analysis of his first-draft design proposal for a distribution system that incorporates 20 ECUs. The system's initial state is now known, and he notices that one particular ECU is used only 10% of the time. Might it be worthwhile to eliminate that ECU and consolidate its functions into one or more of the other modules? Figure 3 illustrates the concept, which is made feasible by the SWC-based AUTOSAR architecture. The engineer manipulates the architecture and reassigns the functions, splitting them into two existing ECUs that have spare capacity.

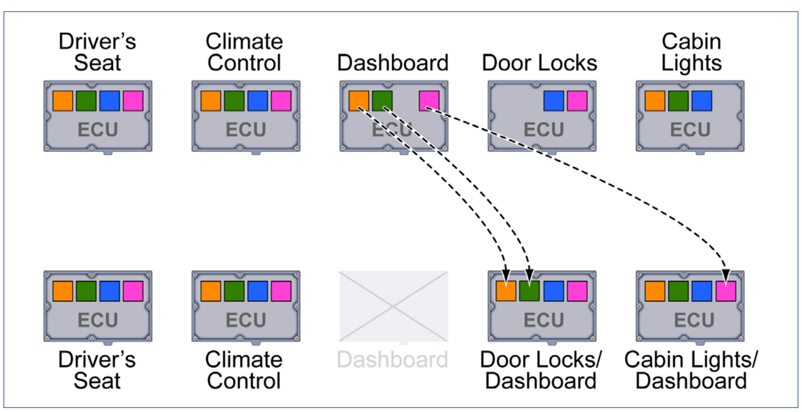

Figure 3: Optimization allocates resources and weighs the consequences, repeating the operation until the best results are achieved. In this conceptual view, five ECUs become four after the SWCs in three modules are consolidated within two ECUs. Side-effects such as cabling and cost impacts must be considered as part of the optimization.

The reassignment presumably will work, since the functional demands on the eliminated ECU were relatively light. But suppose the changes require additional cabling? This might be prohibitive in cost, or weight, or routing complexity. As explained earlier, the AUTOSAR analysis tools should know everything there is to know about the system characteristics. The analysis detects conflicts with parameters according to a hierarchy of priorities. Figure 4 depicts a typical graph plotting key metrics in an optimization hieraracy. Perhaps cost is the highest priority and the ECU change adds more cabling cost than it saves in ECU costs. While the ultimate decision is still left to the engineer, at least he has some concrete input from which to make his choices. He assesses the results from as many manipulation/analysis cycles as time permits, and the best outcome—as judged by the designer himself—is an "optimized design."

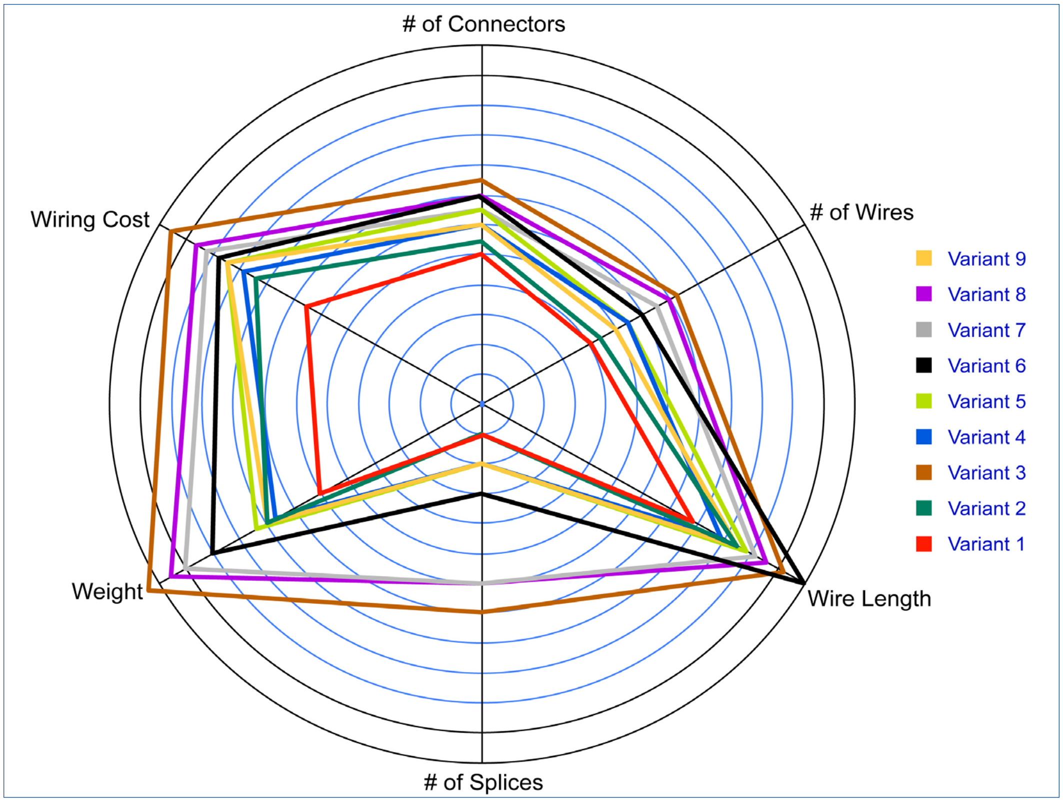

Figure 4: Metrics plot. This circular chart maps six variables in the wiring harness as the engineer evaluates design alternatives.

There are some innate weaknesses in this methodology. How many iterations can the engineer run within, say, a one-month deadline? Two? Five? Each cycle requires time for assessment and time for the next round of manipulation. Assuming for the sake of discussion that the engineer can do 3 full iterations in 30 days, how does this compare with the total number of combinations possible with 20 modules, hundreds of SWCs, and thousands of other variables?

And of course, there is always a concern about the engineer's preconceptions. Even a very open-minded designer is likely to have some beliefs about what can or cannot be done with the system. In effect these can set arbitrary limits on the testing of new ideas.

These issues point toward the need for some level of automated optimization with AUTOSAR. The underlying analysis technologies might be the same as those used for manual optimization but putting computer-based tools in charge of the manipulation and solution assessment vastly expands the number of alternatives that can be evaluated.

…AND TOMORROW

With AUTOSAR, design automation tools—particularly modeling and related functions—become a practical part of the vehicle engineer's kit. Semiconductor designers and printed circuit board developers have long relied on automated optimization to advance the state of their art. Commercially available integrated circuit (IC) optimizers such as Olympus from Mentor Graphics have helped IC gate counts grow and feature sizes shrink by multiple orders of magnitude over a twenty-year span.

Today's reality for automotive engineers remains manual optimization. AUTOSAR has given designers the means to describe a system in a structured way and to use the stored knowledge to calculate the effects of changes. It is a powerful first step that relies on the designer's instincts and experience.

But the growing complexity described at the beginning of this article is bringing pressure to extend the computer's influence in optimization activities. Clearly the next step forward is non-intrusive guidance that helps the designer navigate the terrain between design alternatives. Guided optimization can direct the user in a sequence of logical moves toward an optimized result, and can help overcome assumptions that might hamper exploration. Guided optimization provides a human operator with easier ways to manipulate the system and monitor the resulting state changes. Solutions for this discipline are emerging now, with individual tools aimed at architecture design and metrics, bus simulation, and more.

Extrapolating the trend line yet further, fully automated optimization with AUTOSAR is the advancement that promises to pay back all the effort that must go into developing the technology. With true automation, the computer can attempt the millions and even billions of trials that are required to evaluate all possible combinations of elements (SWCs, wiring, cost, etc.) and behaviors. It can quickly narrow the range of viable choices to a human-comprehensible scale. It can produce reports that estimate total cost savings per vehicle. And equally important, it can complete these operations in minutes or hours.

But fully automated optimization has some critical barriers to cross before it can be widely accepted or adopted:

- AUTOSAR itself must achieve broad acceptance among auto makers around the world.
- The first deliverable vehicles incorporating the standard are only now coming off the production lines, and just a fraction of these cars' subsystems are AUTOSAR-compliant. The greatest benefit of optimization will come when vehicles are compliant from bumper to bumper.
- Optimization technology itself is still on the drawing boards. While automatic optimization in general is relatively well understood, a toolset that can deal with today's complex automotive systems is itself a complicated design challenge. Moreover, the kind of uniform databases that can support optimization are not in place.
- Broad market acceptance of either AUTOSAR or its optimization tools is not guaranteed. Risk-averse OEMs may be slow to adopt new, unproven methods.

Nevertheless, there are powerful forces pushing the market in the direction of increased standardization and the design automation it enables. Chief among these, as always, is cost. In a recent "real-life" example, a major manufacturer of prestige passenger cars saved tens of millions of dollars on its new model after an optimization effort revealed a way to eliminate a single ECU from the vehicle's design. While this was not an automated optimization, its success still highlights the importance of finding and addressing unnecessary costs. Even a minimal amount of optimization can have an impact on profitability and competitiveness.

CONCLUSION

Standards are the enabling platform for the modern computer-based design tools that have transformed industries around the world. AUTOSAR is a leading effort to bring some standardization to the software platforms that are at the heart of today's automobiles. AUTOSAR is gaining ever wider acceptance among car makers, which sets the stage for further development of automated tools for system design, evaluation, and feasibility studies. As these tools begin to penetrate the automotive design field, they will become the foundation of optimization processes that help engineers design profitable automotive products more efficiently than ever before.

Serge Leef, General Manager, System Level Engineering Division
serge_leef@mentor.com

Please direct any questions or enquiries to

James Price, Automotive Marketing Manager
+44 1 503 310 1652
james_price@mentor.com

or contact your local sales representative
www.mentor.com/automotive

Corporate Headquarters
Mentor Graphics Corporation
8005 SW Boeckman Road
Wilsonville, OR 97070-7777 USA
Phone: 503.685.7000
Fax: 503.685.1204

Sales and Product Information
Phone: 800.547.3000

Silicon Valley
Mentor Graphics Corporation
1001 Ridder Park Drive
San Jose, California 95131 USA
Phone: 408.436.1500
Fax: 408.436.1501

North American Support Center
Phone: 800.547.4303

Europe
Mentor Graphics Deutschland GmbH
Arnulfstrasse 201
80634 Munich
Germany
Phone: +49.89.57096.0
Fax: +49.89.57096.400

Pacific Rim
Mentor Graphics (Taiwan)
Room 1001, 10F
International Trade Building
No. 333, Section 1, Keelung Road
Taipei, Taiwan, ROC
Phone: 886.2.87252000
Fax: 886.2.27576027

Japan
Mentor Graphics Japan Co., Ltd.
Gotenyama Garden
7-35, Kita-Shinagawa 4-chome
Shinagawa-Ku, Tokyo
Japan 140-0001
Phone: +81.3.5488.3030
Fax: +81.3.5488.3021

MGC 11-09 TECH8760-W