PARALLEL PERFORMANCE OF LS-DYNA ON THE NEXT GENERATION OF COMPUTER SYSTEMS

Jeff Zais, IBM

zais@us.ibm.com
715-386-6424

Abstract

LS-DYNA is available on a wide variety of computer platforms, ranging from commonplace personal computers to sophisticated vector processors to high-performance scalable servers. This presentation will review historical LS-DYNA performance and look forward to what can be expected in the next generation of computer platforms.

The capability of available computer platforms always has a bearing on what type of simulations LS-DYNA end users can run. In the past, departmental machines such as DEC VAX servers provided many computing cycles to users. Porting of the code to Cray vector machines made overnight full-car crash simulations practical. Models grew larger but could still be run overnight as machine speeds increased and SMP parallel versions became available. At the same time, UNIX workstations increased in capability to the point where many sophisticated simulations could be run on desktop machines.

Most recently, significant advances in LS-DYNA and computer hardware have taken place in two different areas: scalable servers, and machines based on low-cost processors. The advances are being driven by the growing acceptance of the MPI version of LS-DYNA, the level of performance available from low-cost commodity processors, and the level of performance available from high-performance processors available in scalable systems. Current performance examples and trends in microprocessor improvement can be used to give users guidance on what levels of LS-DYNA performance will be attainable in the next few years.

Performance of LS-DYNA in the Past Decades

This year marks the 25th anniversary of LS-DYNA and its predecessor, DYNA3D. Since the first release of LS-DYNA in 1976, performance of LS-DYNA has increased greatly due to several factors. This paper looks back at the major contributors to this increased level of performance, and also looks forward at what LS-DYNA users can expect in the next decade.

Vector Machines

DYNA3D was first released in 1976. The code has traditionally been available on a wide variety of computer platforms. Optimization and tuning of LS-DYNA on these platforms has greatly benefited the end users. When DYNA3D was initially released, DEC VAX machines were typically used for scientific applications. This provided an adequate level of performance for DYNA3D users, but a major advance was the first CRAY-1 vectorized version, available in 1979. A decade later, at the end of the 1980's, shell element formulations were fully vectorized on the CRAY Y-MP, leading to performance improvements of up to an order of magnitude. This in turn led to the popularity of LS-DYNA in the automotive industry. At about the same time, in 1989, a version for the IBM PC was also first made available.

For the decade 1985-1995 vector machines were the clear leader in LS-DYNA performance, running simulations far faster than the available RISC (reduced instruction set computing) platforms. Not only was the single CPU performance far superior, but SMP (shared memory programming) parallel processing was also first developed on the CRAY X-MP and CRAY Y-MP vector machines.

Emergence of RISC Machines

The 1990's began with RISC machines more than an order of magnitude behind vector machines in performance. A significant change was observed with the introduction of the IBM RS/6000 590 workstation, based on the POWER2 processor. With this machine, RISC performance took a significant step forward, and RISC performance in the 1990's has improved continuously as processor speed has increased. A good indication of recent progress is shown in this set of data on elapsed time performance for a Volvo metalforming simulation

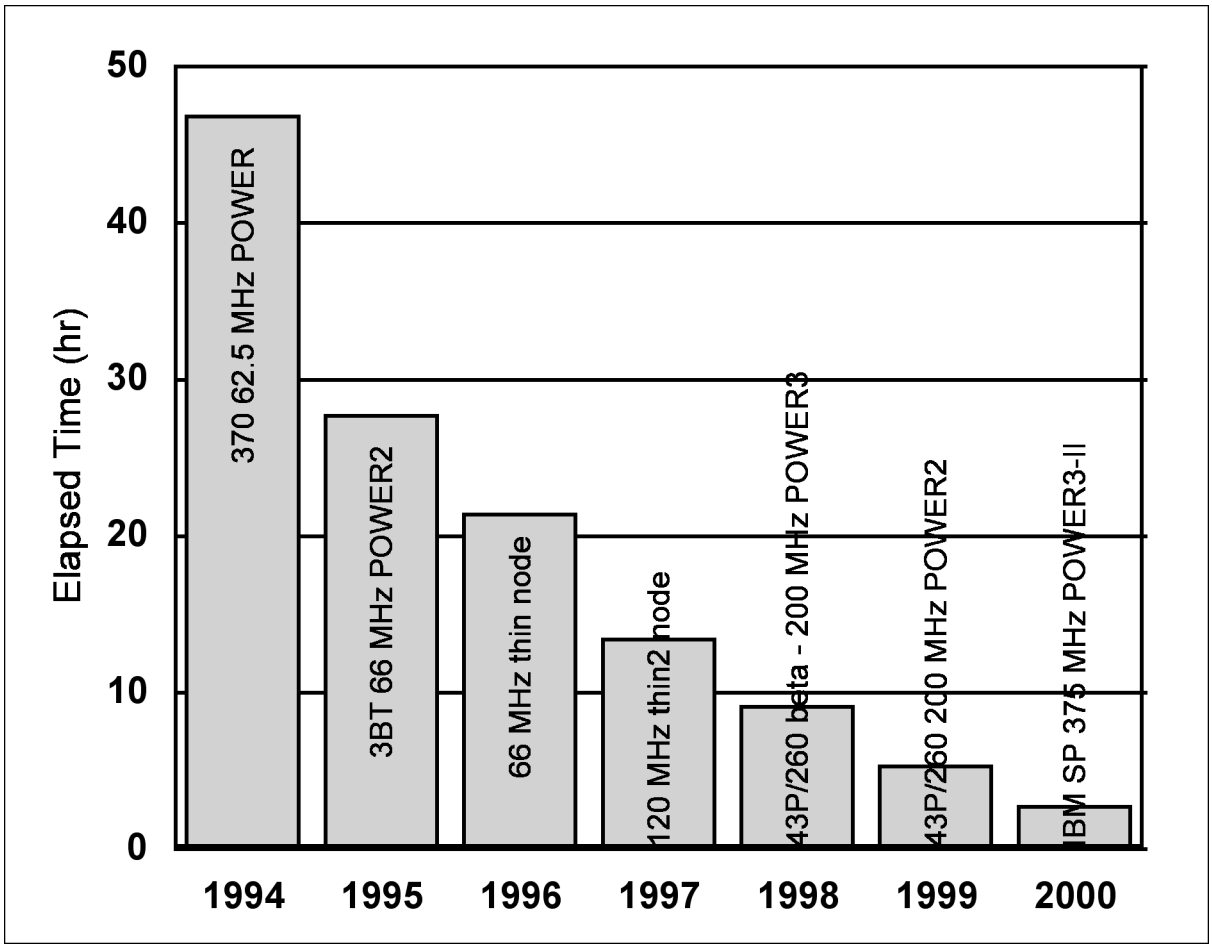

Figure 1. Elapsed time over past seven years, on a single CPU, of Volvo's 71,000 element NUMISHEET93 stamping simulation, using the fastest available IBM processor.

This data shows that a simulation which originally took 2 days can now be run in less than 3 hours, a gain of 17.3 (computed from the elapsed time ratio 46.8 hours/2.7 hours).

RISC CPU single processor performance was still well below leading vector machine performance during this period, but even for high-end simulations, two factors led to the popularity of RISC platforms:

1. superior price/performance; and

2. SMP parallelism allowed for relatively high performance.

Emergence of SMP Parallelism

During this same period, the SMP parallel capability pioneered on vector machines was transferred to the RISC machines. A key factor in customer acceptability was the development of the consistent results option, popularized in versions 940 and 950 of LS-DYNA. Given confidence in parallel results, customers rapidly accepted the SMP parallel version, which provided them a substantial boost in performance, typically a factor of 3 when using 4 processors. In the latter half of the 1990's, using SMP parallelism with 4 or 6 processors, a RISC server could attain about half the performance of the fastest available vector machines. This high level of performance proved very attractive to all but the customers who required the fastest turnaround times, who chose to continue with vector solutions until MPI parallelism became widely accepted.

Emergence of MPI Parallelism

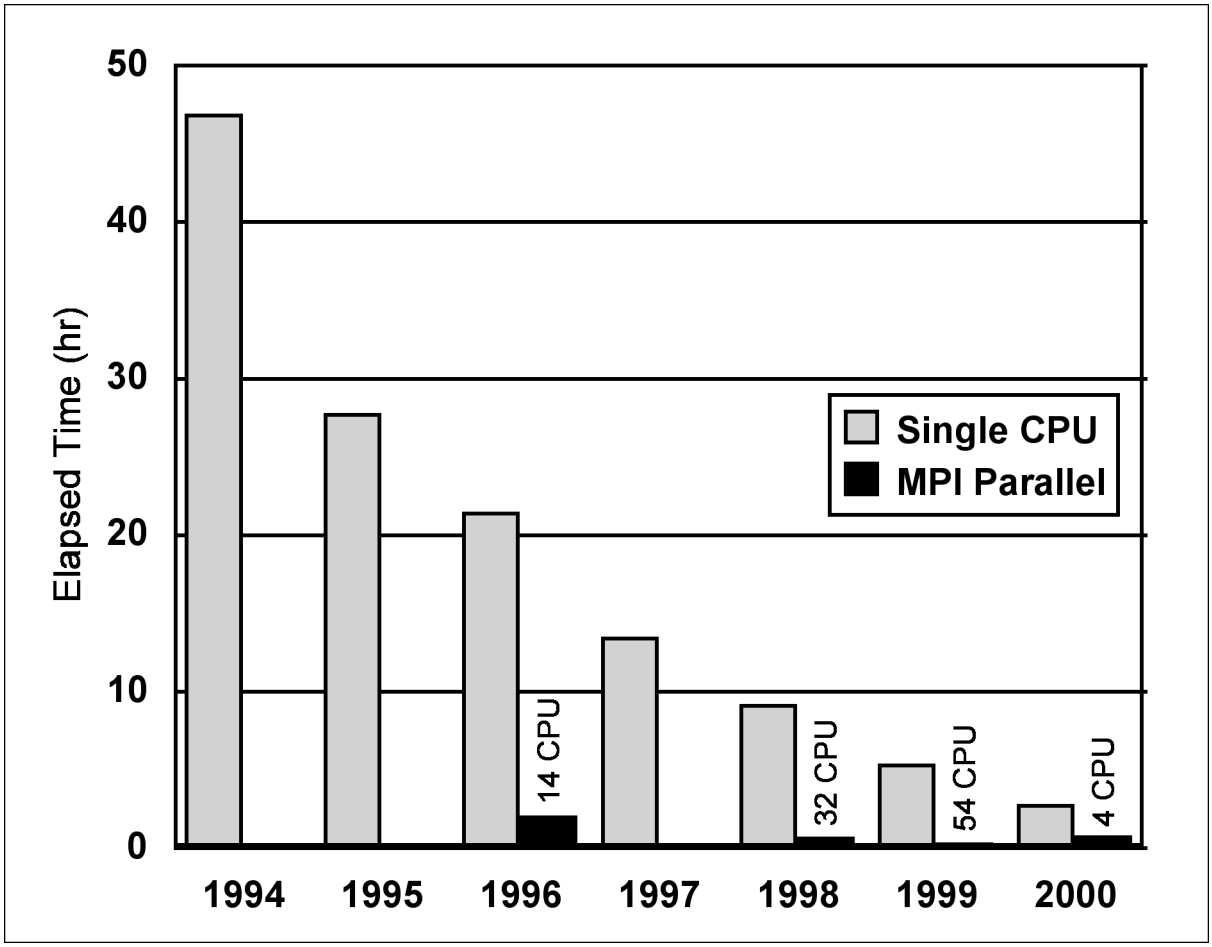

Development of the MPI parallel version at LSTC (often known as MPP-DYNA) started in the early 1990's and progressed throughout the decade, culminating with several trials at customer sites. One of the first sites where this code was put into production was Volvo's Olofstrom location, where metalforming simulations were performed in parallel on an IBM SP. The MPI results can be added to the data from Figure 1 to show even more dramatic advances in performance

Figure 2. Elapsed time of Volvo's 71000 NUMISHEET93 stamping simulation run on single and multiple CPUs, using fastest available IBM processor and various configurations of IBM SP systems.

This data shows that the job's elapsed time in 1999 had been reduced from 5.3 hours to 16 minutes through the use of MPI parallelism. Today, MPI parallelism on a large car crash model can produce typical parallel speedups of 20 on 32 processors, or 12 on 16 processors.

Performance of LS-DYNA Today

The widespread acceptance of the MPI version of LS-DYNA has completed the shift away from single powerful processors to a cluster or network of processors. These processors can be tightly coupled in the same machine, or linked together with a communication network. For LS-DYNA today, both the processors and the networks are important factors in determining the level of performance.

Emergence of Commodity Processors and Proprietary RISC Processors

As pointed out in the previous sections, LS-DYNA has historically been available on a wide variety of platforms, and first appeared on an IBM PC version in 1979. Since that time, processor technology has rapidly advanced, to the point where today's commodity processors can perform LS-DYNA simulations at nearly the same rate as RISC processors. Two comparisons demonstrating this are shown in this table of elapsed time for single processor runs on two publicly available input decks:

Model	*Pentium-III 600 MHz*	*Pentium 4 1500 MHz*	*POWER3-II 375 MHz*
NCAC Neon model 272,180 elements 80 msec simulation	92.7	46.3	44.5
NCAC C2500 model 54,028 elements 80 msec simulation	<no data>	9.34	8.11

Table 1. Elapsed time (in hours) of two typical LS-DYNA input decks run on a single processor. Pentium times from IBM xSeries 330 servers. POWER3-II times from an IBM SP system.

The plans disclosed for the IBM Power4 processor designed with a clock speed in excess of 1000 MHz could mean a large increase in performance, so that for the next few years, a two-tier performance system could be in place - commodity chips (Pentium 4 and follow-ons) at one level of performance, with proprietary chips (such as the Power4) at about twice that level.

Importance of Communication

As usage of the MPI version of LS-DYNA on computing clusters becomes more widespread, the communication network becomes more important. Communication over the network is not as fast as within a node, but communication technology is also advancing at a steady rate. The important characteristics of a network are the latency (time to get a piece of data from a remote location) and the bandwidth (sustained speed at which data can be transferred between nodes in the cluster). Some representative values for different types of networks are:

Communication Characteristic	*within an IBM SP node*	*IBM SP Switch*	*IBM SP Switch2*	*100T Ethernet*	*Myrinet*
Latency (micro sec)	7	20	17	85	18
Memory Bandwidth (MB/sec)	394	136	330	11	140

Table 2. Characteristics of some representative communication networks used to connect computational nodes.

Just as processors have evolved into two levels, high-performance processors and commodity processors with a somewhat lower level of performance, communication networks also range from high-performance to low-cost options. The effect of these networks on LS-DYNA performance can be demonstrated with the some examples:

The following data is for a 20 msec simulation of a 300,000+ element model. One set was run on an IBM SP with a high-performance IBM Switch2, and another was run with a somewhat slower gigabit Ethernet.

Processors	*IBM SP with gigabit Ethernet*	*IBM SP with Switch2*
8	2.12 (1 x 8 CPU)	2.12 (1 x 8 CPU)
16	1.18 (1 x 16 CPU) 1.28 (2 x 8 CPU)	1.18 (1 x 16 CPU) 1.19 (2 x 8 CPU)
24	<no data>	0.85 (3 x 8 CPU)
32	1.02 (2 x 16 CPU)	0.66 (4 x 8 CPU)
48	<no data>	0.52 (4 x 12 CPU)
64	2.52 (4 x 16 CPU)	0.49 (4 x 16 CPU)

Table 3. Elapsed time (in hours) for 20 msec simulation of large 300,000+ element model. For a high number of processors (32 and above) the high-performance SP Switch2 is required for scalability.

This demonstrates that for the highest-performing networks, a model of this size can scale effectively to 48 or 64 processors. A network such as gigabit Ethernet does not allow scaling much beyond the single node with 16 CPU.

A high speed network like the IBM SP Switch enables the user to run MPI LS-DYNA to many processors. As a test of this capability, the NCAC Caravan model was run on an IBM SP with 4 CPU Winterhawk-2 nodes. The elapsed times for this experiment show that the practical limit to scalability is about 64 processors, but that speedups are observed all the way to 224 CPU before overhead from the parallelism causes an increase in elapsed time.

Processors	1	16	32	48	64	96	120	192	224	238
Elapsed Time (hr)	69.16	6.03	3.3	2.4	1.97	1.66	1.58	1.49	1.48	1.49

Table 4. Elapsed time (in hours) on an IBM SP with POWER3-II processors. High performance SP Switch allows for good scalability to 64 processors, and elapsed time continues to drop as processors are increased all the way to 224.

Data is also available for commodity processors in nodes with fewer processors. In this test, a 10 msec simulation of a large model (400,000+ elements) was run on a cluster of IBM xSeries X330 nodes, each with two processors. The clusters was connected with 100T Ethernet and Myrinet, so both options were tested. The elapsed time for this simulation are:

Processors	600 MHz Pentium III IBM X330 + Myrinet	866 MHz Pentium III IBM X330 + Ethernet	375 MHz POWER3-II IBM SP + SP Switch
4	8.87	5.79	4.15
8	4.28	2.77	1.82
16	2.39	<no data>	0.9

Table 5. Elapsed time (in hours) of a large LS-DYNA input decks. Pentium times from IBM xSeries 330 cluster. POWER3-II times from an IBM SP system.

This indicates that for a moderate number (up to 16) of these commodity processors, good performance can be obtained using the available network options. However, the scaling is not as good as using the high performance SP Switch available in the IBM SP system.

LS-DYNA Performance in the Next 10 Years

Anticipated Advances in Computer Hardware

At a 2000 IEEE Electron Devices conference, Intel reported that the company was able to produce transistors with elements as narrow as eight angstroms, or three atoms, far smaller than today's chips. This led to a prediction that Moore's Law will hold for another decade, with no known physical barriers to slow the pace of chip development.

Moore's Law, formulated by Intel cofounder Gordon Moore, states that the number of transistors on a chip will double every 18 to 24 months, as transistor size shrinks. More transistors lead to a corresponding increase in performance. Early in the 1990's, Moore himself predicted the industry would hit a wall when transistors shrunk to about 0.25 microns. However, this prediction proved too pessimistic, since chips with that size were introduced in 1997.

IBM Microelectronics is already producing 0.13-micron chips, using three advanced techniques developed at IBM: copper wiring, silicon-on-insulator transistors, and a new type of insulating material known as a low-k dielectric. The industry expects moving to 0.07-micron chips in 2005.

Using the previously presented Volvo data, LS-DYNA performance can be compared to observe increases expected from Moore's Law. In this graph, actual increases in performance are compared to doubling of performance every 24 months and every 18 months.

Figure 3. Actual single processor performance observed over the past seven years compared with Moore's law projection of doubling performance every 18 months.

The chart indicates that the actual performance gains over the past 7 years closely follow the projections doubling performance every 18 months.

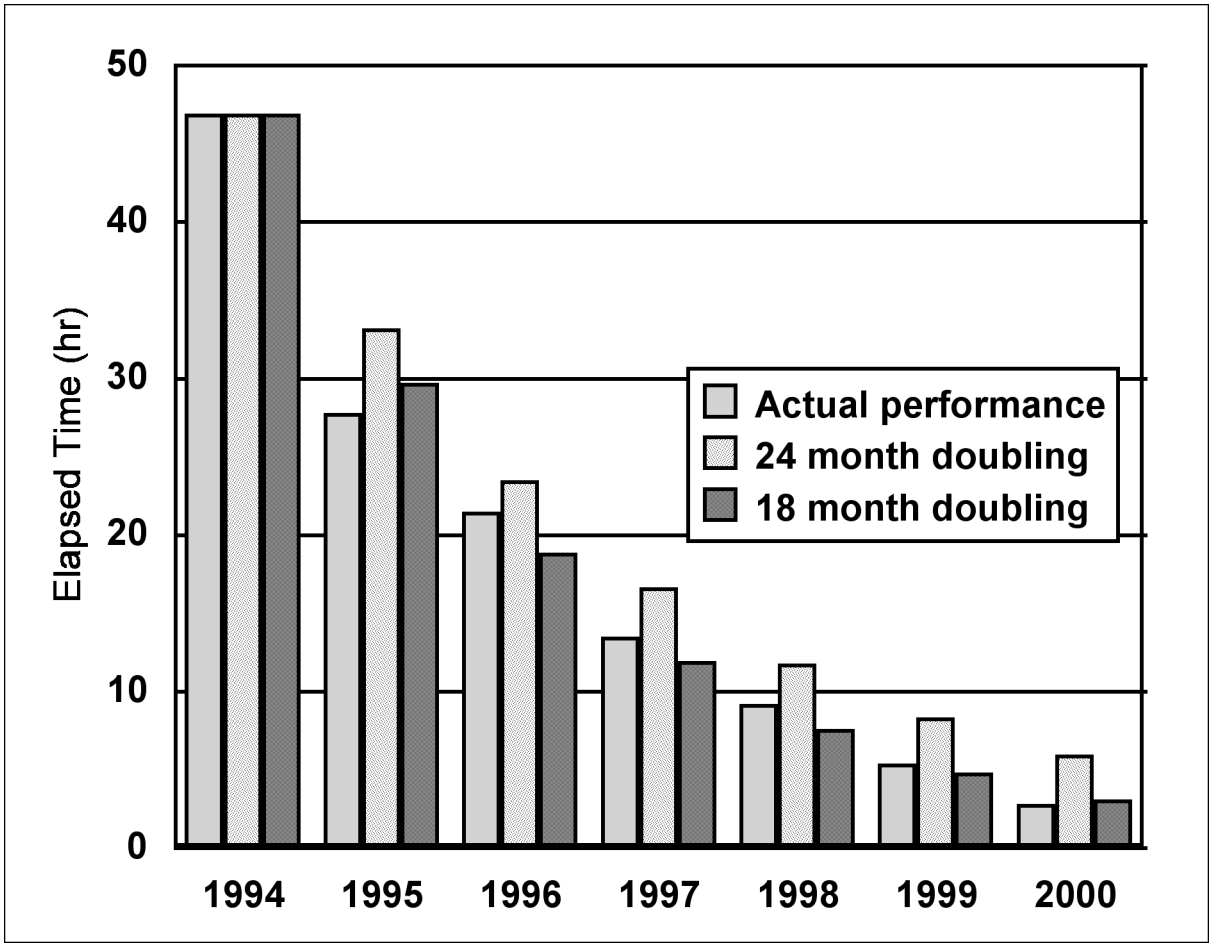

Networks for clusters have been observed to increase capability over the years as well. For instance, the natural progression for Ethernet (10T, 100T, gigabit) is certain to continue so that speed increases for networks match the speed increases in processors.

Anticipated Advances in LS-DYNA Software

The period covered in Figure 2 could be considered the golden age of performance increases in LS-DYNA. In addition to the increases in processor speed, this is the period in which parallelism, first SMP and then MPI, were widely developed and became accepted in the user community. SMP parallelism made possible a performance increase of 3-5 over single processor results. MPI parallelism made possible a performance increase of up to 20 over single processor results.

Unfortunately, no such dramatic performance increases from further parallelism and code optimization are anticipated. In fact, overall performance and parallel scaling of the software could decrease slightly over the next few years. Most of the LS-DYNA development activity will likely be focused on adding additional features and capabilities into the code. A natural side effect is that new features sometimes are not fully optimized and tuned for performance when originally introduced into a code.

Performance Issues over the Next 10 Years

Projecting out to the end of the decade, chips with the eight angstrom transistors will contain 400 million transistors and run at 10 GHz, and also run on less than one volt of power.

This compares to today's Pentium 4 chips, which run at 1.5 GHz and contain 42 million transistors.

How will this affect LS-DYNA performance? Using the more conservative estimate of doubling LS-DYNA performance every 24 months, performance will increase by a factor of 16 over the next 8 years. This would mean that running a 100 msec full-duration 500,000 element crash simulation would require:

	4 CPU	*8 CPU*	*16 CPU*	*32 CPU*
Current technology	48	21	11	7
Increase in performance by factor of 16	6	2.7	1.3	0.9

Table 6. Elapsed time (in hours) for a 100 msec run of a 500,000 element model.

Such performance possibilities will clearly affect what users are able to do with LS-DYNA in the future. Some potential outcomes from increase performance include

larger models - model size has continually grown over the past decades, and this trend will likely continue, though at perhaps a slower pace. The difficulties of both building and post-processing these large models also serve as impediments to model size.

more comprehensive simulations - as the cost of processing continues to drop, several variations of a design can be simultaneously attempted. Some level of optimization, using LS-OPT or some other technique, will be practical for large models.

more accurate simulations - explicit finite element simulations have only been practical because many assumptions are made in order to complete the calculations in a reasonable time. In the past, the additional processing power has generally been applied to faster turnaround and larger models. In the future, users may well apply available computing power towards more accurate solutions.

Acknowledgments

Many thanks are due to Guangye Li of IBM for providing several of the benchmark results. Thanks also to Volvo Olofstrom for the performance figures of the NUMISHEET stamping model. Finally, thanks to the entire staff and development team at LSTC for their assistance.