Final CTI Swiss-Tx Project Report

(June 2001)

Swiss-T1

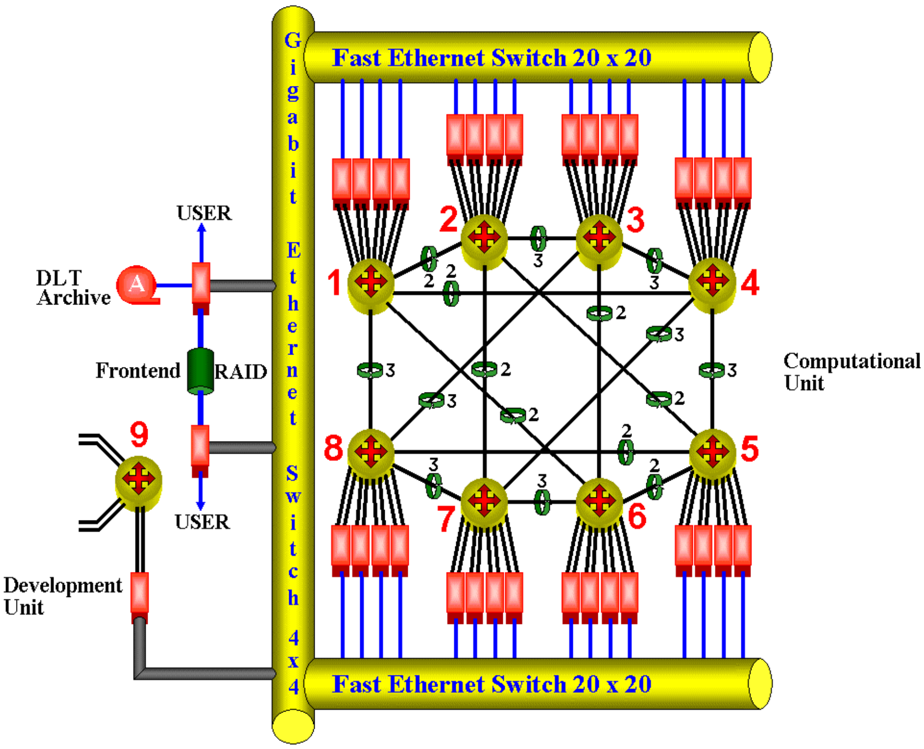

Final CTI Swiss-Tx Project Report

This report compiles all the work performed during the whole Swiss-Tx project that had altogether 4 periods (in parenthesis are the CTI contributions):
Phase 1: June 1st 1998-July 31st 1999 (CTI 3913.1, CHF 780'112.-)
Intermediate phase: August 1st 1999-October 31st 1999 (CTI 4498.1 KTS, CHF 179'137.-)
EPFL sponsored phase over CHF 128'681.- : November 1st 1999-December 31st 1999
Phase 2: January 1st 2000-September 30th 2000 (CTI 4732.1 SUS, CHF 579'066.-)

Preface

The major outcome of the project is the Swiss-T1 cluster computer (see front page) that is presently running in production mode at the EPFL computing centre. Results obtained on the most recent high performance parallel NUMA and MPP machines show that the Swiss-T1 cluster is a highly competitive supercomputer. The Start-up Company ClusterSolutions (CS) S.A. has been created in November 2000 and received the CTI Start-up Label in March 2001. Its business model is to design, build and market application-dedicated cluster computers for engineering, and bioinformatics. This implies that the purpose of the Swiss-Tx project has been fully achieved.

Content

Executive Summary

The Swiss-Tx CTI project has successfully finished at the end of September 2000. The project achieved the goal to build a Swiss-made parallel commodity cluster computer that can productively run in a computing centre. In fact, the Swiss-T1 machine, shown on the cover of this report, has been accepted by the EPFL on July 2000 and was inaugurated on August 23rd 2000. Since then it runs productively. From July to December 2000, the first users took 40% of the maximal possible available time. During the first 3 months in 2001, Swiss-T1's load went up to 70%.

As a consequence of this project, the EPFL Start-up company ClusterSolutions (CS) S.A. has been created on November 2000 and already has important clients. The business plan is added to this report. In addition, Supercomputing Systems A.G. in Zurich (developer of the communication system TNet) was able to achieve a high level of know-how and found an investor in their new technology. This guarantees a further development and a future support of the already installed TNet communication hardware and software. Presently, EPFL considers an upgrade of the Swiss-T1 machine in 2001.

In fact, the project lasted for two and a half years. It benefited from collaborative work coming from academia and industry. It has been one of the rare projects where EPFL, ETHZ and CSCS worked successfully together. Besides these valuable partners, Compaq and Supercomputing Systems joined the project as industrial partners. The project profited a lot from a close co-operation with American research institutions such as Sandia National Laboratory and Oak Ridge National Laboratory. Altogether 4 prototype machines have been built, the last two, the Baby T1 and the Swiss-T1, became production machines that run successfully at the EPFL, and the Swiss-T0 will be reused in another CTI project.

We can claim that the Swiss-Tx project fully satisfied its objective. The cluster technology has been designed, built and fully validated. This technology is now very rapidly gaining importance on the world market.

1 From supercomputers to commodity cluster computers

The very strong evolution in the supercomputer market towards cluster machines that we predicted in the initial Swiss-Tx proposal (http://capawww.epfl.ch/swiss-tx) indeed happened. The recently most powerful machines were ordered at Compaq by CEA (5 Teraflop/s), Pittsburgh University (6 Tflop/s) and LLNL (10 Teraflop/s). As a consequence, the traditional vector supercomputers are slowly disappearing from the HPCN market, NEC and Cray remaining the only vendors of vector machines.

The now popular SMP and NUMA servers use commodity processors, but still have many customised elements. The OS, the compilers, and the communication network are vendor dependent and the methodology to get the power out of the parallel machine has to be adapted to the special underlying architecture. If a user needs a machine with more performance than the one offered by the biggest NUMA machine, a clustering of the NUMA machines is needed. This demands an additional strong adaptation of the user programs to the chosen network solution by introducing the MPI library calls for data transfer between the clusters.

In the Swiss-Tx project the commodity components delivered by Compaq are connected by the communication system TNet, specially developed by Supercomputing Systems AG (SCS) in Zürich. The Swiss-T1 machine has been installed at EPFL in January 2000, was accepted in July 2000, and runs now successfully in production mode at the EPFL computing centre. One of the major advantages of the Swiss-Tx approach is its high flexibility that enables it to adapt the network topology to the needs of the user applications. In fact, two communication network systems, the EasyNet bus and the TNet switching network, were built. There is one single communication library, MPI, efficiently implemented on top of the hardware implemented Fast Communication Interface (FCI). MPI is the standard communication library used in user programs that run efficiently on all parallel machines, also on vector machines, SMPs and NUMAs.

During the Swiss-Tx project, several machines have been built, installed, tested and productively used by the user community. These machines and their networks are (in parenthesis are installation dates):

Swiss-T0 (December 1997): 8 Alpha 21164 processors, 500 MHz, 2 GB memory, 64 GB disk, EasyNet bus
Swiss-T0 (Dual) (October 1998): 8 dual processors Alpha 21164 servers, 533 MHz, 4 GB memory, 170 GB disk, EasyNet
Baby T1@DGM[1] (August 1999): 8 dual processors Alpha 21264 servers, 500 MHz, 9 GB memory, 170 GB disk, TNet
Swiss-T1@SIC[2]: 35 dual processors Alpha 21264 servers, 500 MHz, 37 GB memory, 900 GB disk, TNet

The Digital Unix and Compaq Tru64 Unix were used on all clusters, Windows NT was successfully installed on the T0 (Dual) during the period from October to December 1998. The EasyNet bus communication system by SCS was used to test the Remote Store strategy used in their final TNet network.

In parallel supercomputing, both the computing network latency and bandwidth are crucial. One of the goals of supercomputer manufacturers is to minimise the latency and maximise the data bandwidth. The supercomputer manufacturers of the past created customised supercomputers with custom hardware, custom operating systems and proprietary software libraries. The communication systems were adapted to most demanding operating systems and applications. This full-custom system design led to high development efforts (time and money) and to high-price general-purpose machines.

---

[1] DGM = Département de Génie Mécanique de l'EPFL
[2] SIC = Service Informatique Central de l'EPFL

The goals of today's supercomputer manufacturers are the following:

- Use commodity parts as much as possible to reduce development efforts
  - Use standard PCs, workstations and servers instead of specially developed custom computers
  - Use standard operating systems rather than to design custom operating systems or modifying existing operating systems (patches, fixes, adaptations to new hardware)
  - Use standards (PCI, fibre channel, 19" rack elements, MPI)
- Reduce latency and maximise data bandwidth using advanced hardware and software architectures designed for standard communication libraries
  - Usage of zero-copy principle for hardware and software design
  - Usage of DMA capabilities for sending and receiving data
  - Support for MPI at hardware level
  - Usage of a reduced stack layering for communication
- Allow easy management of parallel commodity supercomputers
  - Adaptive design of cluster computer architecture (torus, mesh, K-ring, hypercube, fat tree)
  - Monitoring and servicing of the supercomputer system
  - On-line modification of the running system
  - Debugger and tracer support for application development
  - Remote control of the machine

The goals of the Swiss-Tx supercomputer project were to design and build a supercomputer system from as many commodity parts as possible and reduce the development effort only to the most necessary parts needed for *parallel commodity cluster computing*. Those are:

- Communication hardware to enable fast communication (low latency, high bandwidth)
- Drivers for the communication hardware for the most popular operating systems (LINUX, Compaq Tru64 UNIX (former DIGITAL UNIX), Microsoft Windows NT)
- Communication software libraries for developers (FCI, MPI for C and FORTRAN)
- Debugger and tracer support for development of distributed applications
- Tools for computer aided and remote management of supercomputers
- Parallel I/O library for large distributed data files

All other parts were taken from the open market for computers. These parts are:

- Workstations and servers with powerful Alpha RISC microprocessors from Compaq, large memories, storage devices with large capacity and additional standard computer components (Ethernet, SCSI, VGA)
- Backup, archive and storage devices for large amount of data
- Console switches to select computers for local maintenance and management
- Racks for the storage of all hardware components
- Standard operating systems (LINUX, Compaq Tru64 UNIX and Windows NT) with standard tools and applications (compiler, resource management tools, debuggers, analysers, development environment, system environmental applications, mathematical libraries).

As a consequence, the EPFL Start-up Company ClusterSolutions S.A. (http://www.clustersolutions.com) has been created in November 2000. Its missions are:

- Commercialisation of application dedicated cluster machines in engineering and bioinformatics
- Consulting in cluster architecture and parallel applications

This Start-up Company joined the UDT Group (http://www.udt.ch) that is at the Bern Stock Exchange since March 2001.

2 Results achieved during the Swiss-Tx project

The Swiss-Tx cluster computing project has accomplished the following goals:

- Two supercomputer communication network systems have been built
  - **EasyNet**
    The bus-based PCI adapter «EasyNet», achieving a data bandwidth of 35 MByte/s (hardware performance 48 MByte/s) and a MPI application data latency of 15 µs (FCI application latency for one-sided communications 5 µs). This hardware is limited to 8 computing nodes and was used in the Swiss-T0 and Swiss-T0(Dual) supercomputers.
  - **TNet**
    The Gigabit-Ethernet- and switch-based 32 bit/33 MHz PCI adapter «TNet» reaches a data bandwidth of up to 100 MByte/s in each direction (input and output). The data latency for MPI applications is 15 µs (FCI application latency for one-sided communications 5 µs). The network architecture and size is soft-coded in the network components and can easily be altered.
- Four supercomputers have been built and three of them are still running
  - **Swiss-T0**
    8 workstation computers with one 21164 Alpha RISC CPU 500MHz, 256 MByte RAM and 8 GByte disk storage each. The network used is the «EasyNet» architecture. The operating system is DIGITAL UNIX; the development environment contains FCI, MPI, ScaLAPACK and BLAS for both C++ and FORTRAN compilers. The peak performance of the cluster computer is 8 GFLOPS. This machine will be used in 2001 by the "Tissue cutting" CTI project.
  - **Swiss-T0(Dual)**
    8 DIGITAL AlphaServer 1200 computers with two 21164 Alpha RISC CPUs 533 MHz, 1 GByte RAM and 22 GByte disk storage each. The network used is the «EasyNet» architecture. The operating system was first Windows NT Server Version 4.0 and was later changed to DIGITAL UNIX; the development environment contains FCI, MPI, ScaLAPACK and BLAS for both C++ and FORTRAN compilers. The peak performance of the cluster computer is 17 GFLOPS. This machine was upgraded to the Baby T1 machine.
  - **BabyT1**
    8 Compaq DS20 servers with two 21264 Alpha RISC CPUs 500 MHz, 1 GByte RAM and 22 GByte disk storage each. A Tnet crossbar switch interconnects six servers. The operating system is Compaq Tru64 UNIX (former DIGITAL UNIX); the development environment contains FCI, MPI, ScaLAPACK and BLAS for C, C++ and FORTRAN compilers. The peak performance of the TNet connected part of the Baby T1 is 12 GFLOPS. This machine runs productively since October 1999 and is presently used as the Mechanical Engineering's departmental compute server.
  - **Swiss-T1**
    35 Compaq AlphaServer DS20e computers (32 for the computational nodes, 2 for the frontends, 1 for development purposes) with two 21264 Alpha RISC CPUs 500 MHz, 1 GB RAM (2 GB for frontends) and 22 GByte disk storage each. The network used is a TNet architecture with eight crossbars and 64 PCI adapters to connect the 64 computational processors. The operating system is Compaq Tru64 UNIX; the development environment contains FCI, MPI, ScaLAPACK and BLAS for C, C++ and FORTRAN compilers. The peak performance of the computational nodes of the supercomputer is 64 GFLOPS. The Swiss-T1 machine has been installed at EPFL's computing centre in January 2000. It runs in production mode since August 2000. The usage is about 40% of all the possible CPU seconds, a very good result for this new post-prototype computer architecture (see details in § 3).
- Two supercomputer communication libraries have been built and are running
  - **FCI (Fast Communication Interface)**
    FCI is a versatile communication library allowing quick and easy communication within commodity cluster computers of the Swiss-Tx supercomputer class. It contains only the most necessary procedures for shared-memory-like and message passing applications. The advantages are the reduced latency and high data bandwidth together

with reduced development efforts for application developers. The standardised communication library MPI has been developed on top of FCI. The FCI works on the platforms Compaq Tru64 UNIX, Windows NT and LINUX.

- **MPI (Message Passing Interface)**
  MPI is a standardised supercomputer communication library used on all supercomputers today. Many supercomputer applications are developed for MPI and can therefore be easily ported. The advantage is the high penetration and user acceptance. The disadvantage is the high training period and high implementation effort. The MPI works on the platforms Compaq Tru64 UNIX (former DIGITAL UNIX), Windows NT and LINUX. MPI programs can easily be ported on all the commercialised computing platforms.

- The resource management system and programming tools have been tested and are running
  - **LSF (Load Sharing Facility)**
    LSF is a program to help sharing the computational resources among the different users of the machines. A set of scripts has been written, tested and applied to make best use of the machines.
  - Codine/GRD
    The resource management system Codine/GRD has been installed on the Swiss-T1. Like LSF, it shares the resources of the machine. In fact, user login to the frontend from which the Codine/GRD system submits batch jobs to the computational unit of the Swiss-T1 machine.
  - **Statistics and accounting**
    A user statistics and accounting application is available on the Swiss-Tx machines.
  - **Performance, monitoring and tuning**
    A set of tools is available on Swiss-Tx for those functions.
  - **Totalview (Parallel debugger)**
    Totalview has been installed and tested on the Swiss-T0. It offers the possibility to debug a parallel program. It can presently only be applied to programs using MPICH running over the Fast Ethernet.
- Parallel data handling facilities have been developed and ported to Swiss-Tx
  - **SFIO (High-level parallel file stripping for MPI I/O)**
    A parallel file striping environment has been developed to be used for MPI I/O. All the different MPI data types are covered.
- Validation benchmarks of the Swiss-Tx machines
  - **Parkbench**
    Parkbench measurements have been performed on all the machines.
  - **Applications**
    A set of applications has been ported and executed on the Swiss-Tx machines for benchmark purposes. Performance comparisons with other machines have been made.
- Bring efficient parallel matrix libraries on Swiss-Tx
  - **BLAS, LAPACK, ScaLAPACK**
    Together with Jack Dongarra invited from University of Tennessee at Knoxville, a special effort has been made to get highly optimised standard mathematical library routines on the Swiss-Tx.
- A new activity on topologies and optimal routing tables has been started
  - First results have been applied to the architecture of Swiss-T1 and a proposal made for Swiss-T2
- Runtime support for switch-based supercomputing
  - Debugger support for debugging of applications using MPI/FCI
  - Tracer support for visualisation of data routing
- Set of applications for computer aided supercomputer management
  - Tool for designing the supercomputer system architecture by defining the network topology and size.
  - Support for building the supercomputer (correct link connections, configuration of the supercomputer, system start-up)

    - Software for runtime observation/monitoring (finding resource problems of computers, weak points in applications and network topology) and management (on-line partitioning of the supercomputer, resource management)
    - Software for maintenance and service support (software patch installation, application installation, upgrades for the computing network software)
- Development of new hardware systems (with the necessary software)
    - New TNet PCI adapter with the latest technology supporting the 64 bit PCI interface (at 33 MHz) and a dual link. The data bandwidth is 200 MByte/s for each link and direction, delivering a maximum bandwidth of 400 MByte/s per link
    - New PCI adapter – based on the technology of the TNet PCI adapter – to allow low-cost supercomputer networks without using switches. This network adapter can be used for small supercomputer systems for specific applications
    - New TNet crossbar switch to allow communication with higher data bandwidth.
- Additions to the existing hardware (in development)
    - SMP support for FCI and TNet to allow higher system performance by supporting multiple processes on a SMP computer
    - Hardware-supported broadcast and multicast support
    - Support of the generalised memory support (on-board page table of the TNet PCI adapter for using normal memory instead of contiguous memory)
    - New hardware-based fast synchronisation schemes for TNet and FCI
    - System channel support of the TNet PCI adapter for socket communication
- Development of a library similar to the SHMEM library of CRAY supercomputers to allow easy application migration from CRAY supercomputers to Swiss-T1 supercomputer
- Complete the resource management tool kit
- Tools to optimise the topology of parallel machines
- Tools to optimise the routing between boxes
- Port the MEMCOM data management system on the Swiss-T1
- Run many production applications
- The mathematical library AZTEC for distributed memory computers has been implemented and tested

3 Report on the Swiss-T1 machines

3.1 Its usage

3.1.1 Swiss-T1 Baby

The Baby T1 has been productively used at the EPFL computing centre from October 1999 to August 2000 and then transferred to the Department of Mechanical Engineering where it is productive since September 2000.

The Baby T1 has been productively used at the computing centre during the period January to August 2000 with a rate of 28% of the total possible time. During the rest of the time, the machine was in maintenance, in upgrade status to install and test new software, not fully used or just idle. Since August 2000, the Baby T1 is used as the departmental machine of the Mechanical Engineering Department.

3.1.2 Swiss-T1

The Swiss-T1 with its 64 computational processors has started production in July 2000. During the first 6 months 40% of the total available time has been used (Fig. 1). During the first three months of 2001 (January to March, Fig. 2), its usage has grown to 70% of the total time, thus reaching its limits.

3.1.3 First comparisons

First comparisons have been performed during the Swiss-T1 acceptance phase with different programs [34, 36, 37]. All show a good scaling behaviour of the Swiss-T1 cluster when running communication over the TNet. In an “embarassingly parallel” case [39], the Fast Ethernet communication network was sufficiently powerful.

In January 2001, comparison results between most recent parallel machines have been performed using the LAUTREC program [40, 41] from IRRMA. This program is presently taking 25% of the total available time of the Swiss-T1 machine. In Figure 3 are shown the results obtained in January 2001 on the Swiss-T1 (a cluster with 1000 Mflop/s peak processors), the SGI O3K (a NUMA machine with 800 Mflop/s peak processors), the IBM SP3 (an MPP/NUMA machine with 1500 Mflop/s peak processors), and the Cray T3E (an MPP machine with 1200 Mflop/s peak processors). MPI was used on the T1 and the SP3. On the O3K and on the T3E, the original Shmem optimised LAUTREC code has been run. The results for the T1, the O3K, and the SP3 are very close up to 8 processors, scaling is starting to be lost on the T1 and the SP3 when running on 16 processors. Up to 8 processors, T3E’s speed is about 60% of the other three machines. These measurements clearly show that the new cluster computing approach is as powerful as the NUMA, and MPP supercomputers. The MPP and cluster machines are distributed memory architectures and the interprocessor communication has to be programmed in MPI or in another message passing library such as Shmem. A NUMA machine is a shared memory architecture and the interprocessor communication can be programmed in MPI and Shmem, or can directly be taken care of by the system, with a risk of performance loss.

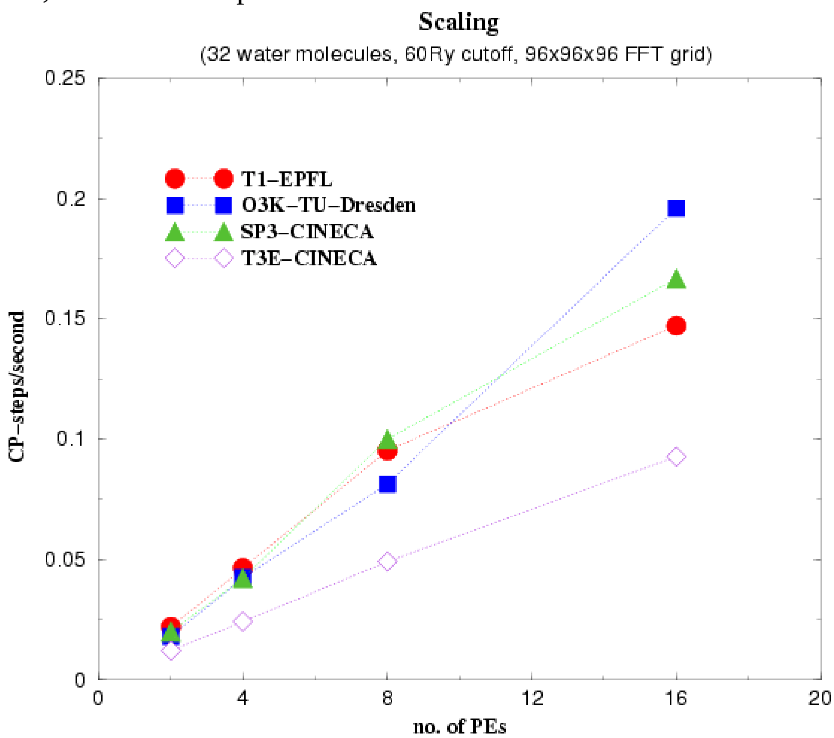

Figure 3 : Scalability study with LAUTREC program on the T1 (cluster, red circles), the SGI O3K (NUMA, blue squares), IBM SP3 (MPP/NUMA, green triangles), and on the Cray T3E (MPP, white quadrangles). Good scaling is found up to 8 processors. T3E is slower than the other parallel machines. MPI has been used for the T1 and the SP3, Shmem for the O3K and the T3E. Four processors are sufficient for this case.

3.2 Expenses for the Swiss-Tx project

CTI contributions

CTI Phase 1 780’112.-

CTI Intermediate Phase	179’137.-	
CTI Phase 2, first 6 months	386’044.-	
CTI Phase 2, last 3 months (under Start-up label condition)	193’022.-	
Total CTI		**1’538‘315.-**
Material and software bought by EPFL		
Swiss-T0	297’247.-	
Swiss T0(Dual)	327’150.-	
Crossbar	80’000.-	
Gigabit Ethernet	40’000.-	
Upgrade T0(Dual) to Baby T1	141’000.-	
Swiss-T1	1’870’424.-	
GRD/Codine	42’000.-	
Total material and software bought by EPFL		**2‘797‘821.-**
EPFL add-ons to CTI contributions		
EPFL 1998	24’205.-	
EPFL 1999	28’795.-	
Salaries Nov/Dec 1999	128’700.-	
EPFL 2000	72’400.-	
Salaries management (estimated)	900’000.-	
Total EPFL add-ons		**1’154‘100.-**
Compaq contributions[3]:		
Add-on to CTI contributions (in cash)	101’207.-	
Engineering salaries (estimated)	150’000.-	
Price reduction for Swiss T0(Dual)	388’500.-	
Price reduction for Upgrade T0(Dual) to Baby T1	167’500.-	
Price reduction for Swiss-T1	811’000.-	
Total Compaq contributions		**1’478’207.-**
SCS direct investment (after deduction of the Swiss-Tx sales profits, estimated by SCS)		**850’000.-**
Total costs of the Swiss-Tx project		**7‘968‘443.-**
		=========

4 The workpackages 1 to 8

The WP1 to 8 (appendices A1 to A8) concern the second phase of the Swiss-Tx project. There is a second WP8 (appendix A9) that corresponds to the WP8 of the first phase. Prof. Maddocks’ group did not contribute to the second phase.

Reports on the workpackages (WP)

WP1: Hardware development (Responsible: Roland Paul)
WP2: Communication software development (Responsible: Martin Frey)
WP3: Resource management (Responsible: Michel Jaunin)
WP4: (Responsible: Roger Hersch)

[3] Compaq made an extraordinary additional price reduction of 38% to the ordinary 30% academic reduction

- Parallel I/O and MPI I/O (Emin Gabrielyan, Roger Hersch)
- HPC in computational structural mechanics (Pieter Volgers)
- NSMB activities in the frame of the Swiss-T1 project (Jan Vos)

WP5: Applications (Responsible: Ralf Gruber)

- The Fluent benchmark results on the Swiss-T1 (Mark Sawley)
- First tests on the Swiss-T1 computer with the AZTEC library (Marzio Sala, Luca Formaggia)
- Installing and benchmarking two webservers: Apache and thttpd (Stefan Schmiedl, Armin Röhrl)
- A MemCom to EnSight native interface toolkit (Jean Favre)

WP6: (Responsible: Pierre Kuonen)

- Computational grid environment and transparent access to Swiss-Tx (Nguyen Tuan Anh)

WP7: Management (Responsible: Jean-Michel Lafourcade)

WP8: Start-up Company ClusterSolutions (Responsible: Jean-Michel Lafourcade)

- Business plan of ClusterSolutions (CS) S.A.

WP8 (1st phase): Data steering and visualisation (Responsible: John Maddocks)

5 Start-up Company ClusterSolutions

November 7, 2000, the EPFL Start-up Company ClusterSolutions S.A. has been created. Its founders are Wenzel Divis, Ralf Gruber, Jean-Michel Lafourcade and Galip Ozem. UDT Group A.G. owns 51% of ClusterSolutions and guarantees financing. UDT Group has entered the stock market in March 2001. ClusterSolutions has got the CTI Start-up Label in March 2001.

ClusterSolutions' missions are consulting in HPTC (High Performance Technical Computing) and design, build and marketing turn-key application dedicated Ready to Run in a Rack cluster solutions in high performance technical computing. The first markets are bioinformatics and engineering machines. First consulting contracts have already been signed.

The financial predictions for 2001, 2002 and 2003 are sales of 5, 15 and 30 Mio. Fr., respectively. It is planned to go for IPO by the end of 2003.

A detailed business plan is added in this report.

6 Conclusions

After the three years period, the Swiss-Tx project has fully satisfied its objective. The particular, highly available, secure, powerful Swiss-Tx cluster technology has been designed, built and fully validated. The Swiss-T1 machine is now in production since July 2000 and runs at 70% of its total maximum possible time. Very recent benchmarks results [41, Fig. 3] have shown that the Swiss-T1 machine is at least still as powerful than the most recent NUMA (IBM SP-3, SGI O3K) and MPP (T3E-1200) machines on the market, and this one year after its installation. Downtime has reached an acceptable low level. What we predicted four years ago has happened; this technology is now very rapidly gaining importance on the world market. The ClusterSolutions (CS) SA Company (see www.ClusterSolutions.com) created in November 2000 has got the CTI Start-up label in March 2001 and markets the Swiss-Tx cluster computing concept.

7 Reports

7.1 CTI related reports

[1] R. Gruber and M.-Ch. Sawley: "Swiss-Tx Series: Project Description", Report for the CTI proposal (CAPA-EPFL, 21.12.97), first phase

[2] Swiss-Tx Team, "Demande de subside", Demande de subside CTI pour 2 ans (CAPA-EPFL, February 1998)

[3] M. Frey, R. Gruber and Swiss-Tx Team, "Swiss-T0 Installation and Acceptance Test Report", Report to accept Swiss-T0 by EPFL (CAPA-EPFL, 27.5.98)

[4] Contract between EPFL and Compaq/Digital, June 11, 1998

[5] R. Gruber and M.-Ch. Sawley, "Swiss-Tx, un Projet d'Entrepreneurs", Flash Informatique (EPFL) FI 6-7 (Juillet 1998), p.1

[6] R. Gruber and M.-Ch. Sawley, "Accord de Collaboration EPFL-Digital", Flash Informatique FI-6-7 (Juillet 1998) p.12

[7] K.M. Decker and R. Gruber, "CSCS/SCSC Contributes to the Swiss-Tx Commodity Parallel Supercomputer Program", CSCS/SCSC CrossCuts 7 (Oct. 1998) Number 2

[8] Agreement between EPFL and SCS AG, December, 14, 1998

[9] N. El Mansouri, R. Gruber and M. Frey, "Swiss-T0(Dual) Installation Test Report", Report on experience made with Windows NT (CAPA-EPFL, 21.12.98)

[10] Memorandum of Understanding (between all the partners), January, 27, 1999

[11] N. El Mansouri and R. Gruber, "Swiss-T0(Dual) Acceptance Test Report", Report to accept Swiss-T0(Dual) under Digital UNIX by EPFL (CAPA-EPFL, 19.2.99)

[12] Swiss-Tx Team, "Intermediate CTI Swiss-Tx Project Report for Phase 1", (CAPA-EPFL, April 1999)

[13] Swiss-Tx Team, "Demande de subside CTI", Prolongation de subside CTI pendant 3 mois (CAPA-EPFL, June 1999)

[14] N. El Mansouri and M. Jaunin, "Baby T1 Acceptance Test Report", (SIC-EPFL, October 1999)

[15] N. El Mansouri et Michel Jaunin, "Résultats de l'évaluation des performances de la machine Swiss-T1", Classified performance comparison report (SIC-EPFL, Octobre 1999)

[16] N. El Mansouri et Michel Jaunin, "Résultats de l'évaluation des réseaux rapides pour clusters Beowulf", Classified performance comparison report (SIC-EPFL, Novembre 1999)

[17] Swiss-Tx group, "Second phase project description of Swiss-Tx Series", Report for the CTI proposal (CAPA-EPFL, November 1999), second phase (1.11.99-31.7.2000)

[18] Swiss-Tx Team, "Demande de subside CTI", Demande de subside CTI 2ème phase (CAPA-EPFL, December 1999)

[19] Business plan for EPFL Start-up ClusterSolutions, December 2000

7.2 Publications

[20] R. Gruber and A. Gunzinger, "The Swiss-Tx Supercomputer Project", Speedup 11 (1997) No 2, p.20-26

[21] R. Gruber and A. Gunzinger, "The Swiss-Tx Supercomputer Project", EPFL Supercomputing Review 9 (1997) p. 21-23

[22] R. Gruber and Swiss-Tx Team: "Swiss-Tx: The Future of Supercomputing?", HPC Profile 20 (1998) issue 20, p.13-14

[23] A. S. Röhrl, "Fast, Portable, Predictable and Scalable Bootstrapping", Proc. of Interface98

[24] J.-A. Ferrez, " Parallel Computation of the Diameter of a Graph on the Cray T3D, the SGI Origin 2000 and the Swiss-T0", in Int. Conf. on Operations Research and annual meeting of GOR, ÖGOR and SIGOPT, ETH Zurich, 31 August - 3 September 1998

[25] J.-A. Ferrez, K. Fukuda and T.M. Liebling, "Parallel Implementation of Graph Diameter Algorithms", EPFL Supercomputing Review 10 (1998) 3-6

[26] R. Gruber and Y. Dubois-Pèlerin, "Swiss-Tx: First experiences on the T0 System", EPFL Supercomputing Review 10 (1998) 19-23

[27] Y. Dubois-Pèlerin, R. Gruber and Swiss-Tx Team, “The Swiss-Tx supercomputer project”, ERCOFTAC journal, 1998
[28] S. Brauss, M. Frey, A. Gunzinger, M. Lienhard and J. Nemecek, “Swiss-Tx Communication Libraries”, Proc. of HPCN99 (Amsterdam, 12-14 April, 1999), 10 pages
[29] P. Kuonen and R. Gruber, “Parallel Computer Architectures for Commodity Computing”, Supercomputing Review 11 (1999) 3-11
[30] Stephan Brauss, “Communication libraries for the Swiss-Tx machines”, EPFL Supercomputing Review 11 (1999) 12-15
[31] B. Gennart, E. Gabrielyan and R.D. Hersch, “Parallel file striping on the Swiss-Tx architecture”, EPFL Supercomputing Review 11 (1999) 15-22
[32] Daniel Weill, “Test de performance du code Speculoos sur l’ordinateur parallèle T0-Dual”, EPFL Supercomputing Review 11 (1999) 29-32
[33] Jean M. Favre, “Visualisation tools and environments for very large data”, EPFL Supercomputing Review 12 (2000) 9-12
[34] M. L. Sawley, T-M. Tran and T.L. Tysinger, “The Fluent 5 benchmarks results on the Swiss-T1”, EPFL Supercomputing Review 12 (2000) 13-16
[35] Emin Gabrielyan, “SFIO, Parallel file striping for MPI-I/O”, EPFL Supercomputing Review 12 (2000) 17-21
[36] Pieter Volgers, “HPC in computational structural mechanics”, EPFL Supercomputing Review 12 (2000) 29-31
[37] Jan B. Vos, “Flow simulations on high performance computers using the NSMB flow solver”, EPFL Supercomputing Review 12 (2000) 32-40
[38] Stephan Brauss, “A new communication architecture for workstation clusters”, Dissertation ETHZ No. 13811 (2000)
[39] M. L. Sawley, P. W. Cleary and J. Ha, “Using smoothed particle hydrodynamics for industrial flow simulation”, EPFL Supercomputing Review 12 (2000) 3-8
[40] A. De Vita, A. Canning, G. Galli, F. Gygi, F. Mauri and R.Car, EPFL Supercomputing Review 6 (1994) 22-26
[41] R. Gruber, P. Volgers, A. De Vita, M. Stengel and the Swiss-Tx Team, “Commodity computing results from the Swiss-Tx project”, submitted to Electronic Notes on FGCS, 2001.