



*Unisys – I gave you all the best years of my life! And it was fun too!*

## PROLOGUE

{This paper by Mr. Loran has been edited for the VIP Club’s Legacy Anthology by LABenson – specifically, most of Mr. Loran’s personal family activities and company parties have been cut out.}

The Unisys Denmark social club UIF (**U**nisys **I**nteresse **F**orening) once had an outing at the Diesel House in Copenhagen - amongst the displays in the museum was a model of the B&W (**BURMEISTER & WAIN**) shipyard buildings. There, I stood by the model and explained about the times that I worked onsite at B&W. After the visit, my boss through many years, Palle Hansen, asked me to write the story of the Unisys sites I have worked at before that history is lost.

Now (2017) after assembling notes for seven years and a 40-year career with Unisys, it is time to finish and publish. Recollections of readers may vary, feedback is welcome, [webmaster@vipclubmn.org](mailto:webmaster@vipclubmn.org)!

### Table of Contents<sup>1</sup>

PROLOGUE.....	1	LUFTHANSA.....	26
INTRODUCTION.....	2	TELIA.....	27
Starting My Story.....	2	AMADEUS.....	29
My Formal Education.....	2	HIS/LMS/DIGISTICS.....	29
Job hunting.....	3	IBERIA.....	30
Name Changes.....	3	TravelSky.....	31
On the Road.....	3	SHORT-TERM CUSTOMER SITES.....	34
Ending My Story.....	3	Fjerndata.....	34
RECKU.....	4	Norsk Hydro.....	34
BURMEISTER & WAIN (B&W).....	6	University of Tampere.....	35
F.L. SMIDTH.....	8	Leitz.....	35
OTHER DANISH SITES.....	11	Natal Building Society.....	35
Dansk Data-service.....	11	Air Canada.....	36
Danish State Railways (DSBundesfahr)..	11	JAT.....	36
DataCentre ved Odense Universitet.....	11	Lloyds TSB.....	36
Synoptik.....	11	ROSEVILLE AND EAGAN VISITS.....	37
Meterologisk Institut.....	12	Courses and Training.....	38
BENCHMARKS.....	12	TDD, UUA/E, USE, NPIT.....	40
DOU Benchmark.....	12	SPECIAL PEOPLE.....	41
Tjæreborg Benchmark.....	13	Tore Ask.....	41
FLS 1100/40-SPU Demo.....	16	Bodil Skovgaard Toftager.....	41
EDB-skolerne Benchmark.....	16	Michael Deneyer.....	42
FLS Cache Disk Demo.....	16	Anders Öhman.....	42
FLS 1100/73 Demo.....	17	Palle Hansen.....	42
Unisys AIDS and COMUS.....	17	Jørn John Loberg.....	42
DBA (COMUS Database):.....	17	Joe Nolan.....	43
MOVING TO ENGLAND.....	18	Ian Hepburn.....	43
AIDS.....	18	Rick Onsgard.....	43
COMUS.....	20	Rob Jamieson.....	44
EAD SUPPORT.....	21	EPILOGUE.....	44
SCANDINAVIAN AIRLINE SYSTEM.....	21	<i>Edited with Microsoft Word, 29-Jul-24</i>	

<sup>1</sup> Click on any section for a quick scroll thereto.

## INTRODUCTION

I decided to organize this story based on the customers (I never grown accustomed to the Unisys term: clients – it has a bad “ring” to it in Danish) that I have worked for, in many countries, with some Unisys-internal memories added here and there. It is organized a little chronologically, starting with my first customer assignment at *Regionale EDB Center ved Københavns Universitet* (RECKU) and then adding customers along the way. However, I will start with some basic background information about myself.

I have inserted ~~many personal stories along the way~~. Maybe it will also tell my wife and children about what has happened during the many days I have been away from home. ~~My wonderful wife, Helle, knows a lot of this, because she has listened to a lot of comments and complaints that I never voiced to the management in Unisys! However, I do not want to cover that in detail here – this paper is primarily a positive story. Without Helle’s support throughout my life; I would never have been able to give Unisys so much nor had so many experiences.~~ *Unisys – you also owe her a lot!*

### Starting My Story

I started in Unisys a few days late (due to a skiing vacation) in January 1974 and since my intended manager Jørn Loberg was not around that day I got “taken” by another manager Svend Baltersen (I think) – and that was my luck, as it meant that I ended up in the) OS1100/2200 world. I will use the name OS2200 throughout, although it was OS1100 for many years. I was immediately sent to RECKU – and that is where my story starts.

### My Formal Education

I have a M.Sc. degree in Electrical Engineering from the Technical University of Denmark in January 1973. There was no Computer Science education in Denmark at that time, so this was probably the more advanced Information Technology (IT) higher education available. I should say that I was not one of the early IT nerds, who spent hours testing out and playing with computers. However, IT was a natural part of the education and I had had several courses involving IT. Funny enough the course in **Operating Systems (OS)** was to a large degree based on MCP – and yet I ended up in the OS2200 side of Unisys.

My master thesis was a program (LACAP) to do numerical analysis of large-scale electrical networks, written in Fortran H for IBM 370 series systems. The University had a “standard” program (ANP3) for this, but it had serious size limitation, and my program expanded the size of the electrical network to be analyzed many times. It required a very tricky, special allocation of temporary mass storage, so already then, I had to search for specialist help to set up the tons of complicated JCL //DD statements correctly. As a curiosum, while I was developing my program, the computer Centre, NEUCC at the Technical University, went through an upgrade from an IBM 360/75 to a 370/165, and in that period, we had access to the RECKU 1108 – my first encounter with a Unisys system. I never made my program work on the 1108 – back then!

My next encounter with computers was during my military service in 1973. I was assigned to a special research group, and in my work, I was asked to write some utilities (backup, file control) for the department’s minicomputer. I did that, but also wrote a golf program that became somewhat popular. It had a built-in facility to checkpoint the golf game when a supervisor walked in, switch to production analysis data, and then be able to restart the game when the boss was gone. I learned a lot about IT during my military time!



### Job hunting

I applied for a few IT-jobs and got one offer before I saw the job advertisement from Univac. I applied and one morning they called and asked me to come for an interview right away [they paid for a taxi all the way in there, unthinkable for a “poor” student – my first encounter with “big business”]. After one interview with H.R. manager Paul Brade, I received an offer for employment [on paper with a masthead of Univac, a division of Sperry Rand]. I therefore started at Unisys, then Sperry Univac, on 1st January 1974. I also applied for a job at Burroughs, but after a few interviews they rejected me. Burroughs later had to buy all of Sperry to just to get me when they realized their mistake – he-he!!! 😊

### Name Changes

The company has changed names many times. In mid-70's a law in Sweden required that government IT contracts should be given to Swedish companies. To win contracts in Sweden Sperry Univac merged with Data-Saab to form Saab Univac, 51% owned by Data-Saab, and this company operated in all the Nordic countries. This gave quite a few large IT-contracts in Sweden. Some years later this law was rescinded, and the company again became Sperry Univac. In the 80's the name was changed to Sperry and finally in 1986 Unisys was formed when Burroughs (Michael Blumenthal) bought Sperry (for money borrowed on the positive value of Sperry)! Throughout this document, I will call the company Unisys, even though it had other names at the times that I am writing about.

### On the Road

I have travelled a lot with Unisys. Although some of these trips have been only airports, hotels, Unisys offices, and computer rooms; I have tried to see and enjoy as much as possible at all the places I have been going.

When I started travelling and went back to the hotel there were no emails to check, no SMS-es (text messages) or no remote access to systems. Thus, I had a lot of time on my own, time to explore, time to go out; and I have maintained that habit ever since.

### Ending My Story

Then 42 years later – it was nearly the end! Not because I stopped, but because U.S. management [probably with the help of U.S.-based consultants] in their wisdom decided to close Unisys Denmark [and eight other countries' offices in Europe] – and with that we all got laid off. What the wise guys failed to understand is that many, many people in these smaller countries contributed positively in other parts of the world [not properly remunerated in their home countries, creating even bigger problems for them]. Some management in Europe, managed to establish an arrangement whereby those employees [primarily ClearPath<sup>®</sup> specialists] got employed by local companies with a guaranteed amount of work from Unisys.

In Denmark, Miracle A/S became the company to help us continue supporting Unisys customers worldwide. Luckily, on the 1<sup>st</sup> of April 2016 Jan Ejner Jensen, Per Anders Wicklund, two Swedes and I transferred to Miracle A/S (still with Palle Hansen as our manager) – thereby continuing our Unisys ClearPath work. For me it was only for about one year, for the others a minimum of three years. When my first year had expired (ended up as only nine months as the deal was delayed by Unisys) I was transferred to a semi-retired, per-hour consultant by 1<sup>st</sup> of January 2017, still employed by Miracle A/S, working for Unisys, for as many or as few hours I felt like – a super arrangement for me.

*But now let history start – rewind back to 1974!*

## RECKU

RECKU stands for “Regionale EDB Center ved Københavns Universitet” – or the regional IT Centre at Copenhagen University. This was one of the three University IT Centres in Denmark at that time; the others were NEUCC at the Technical University (IBM) and RECAU at the University in Aarhus (CDC). These Centres were later combined into one unit, called UNI-C. RECKU was located at the H. C. Ørsted institute on Blegdamsvej in Copenhagen, next to Fælledparken, a large green oasis (highly appreciated by visiting support personnel– some people took very long walks during their breaks!)

In 1974 RECKU had a Unisys 1110 2x1 – a rather large system and one of the busiest [we learned later by a problem we hit before anyone else] sites in the Unisys world then. It was primarily used for students doing IT projects and was only open during the day. It had planned PM (Preventive Maintenance) every day from 7AM to 9AM – a very necessary arrangement during those days. There was a large staff of Computer Engineers (CE) with offices in the basement, amongst them: Frank Vendelsøe-Nielsen, Walther Blom [an American, expert in handling the very unreliable 7015 plated-wire primary memory], Ole Nordahl Nielsen, Per Johansson, and Arne Norholt Olsen.

There was a permanent staff of four System Analysts to handle all the dumps – also required in those days. When I joined, we were Claus Parkhøi [in charge of the software group throughout the existence of this large site – and my mentor], American Paul Ward, British Tim Ward and me. Paul was amazing and I remember Jim Babcock once said: “Paul is like a CULL (a cross-reference program; we had no CULL in those days), just better; the CULL tells you where a routine or variable is referenced – Paul also tells you why. RECKU also had their own staff of highly qualified system support people; I remember Niels Jørgen Hansen [later at F.L. Smidth], Steen Jarbøel, and Lene Veje [both later at Unisys and D.S.B.]. RECKU also had some good machine operators [they got lots of practical experience in problem handling in those days] of which I later worked a lot with Jeanne Eismark when she came to SAS.

I was right away sent to OS2200 courses, with the first one being an ECL course by Jan Kennet Larsen [an amazingly skilled teacher – still is!] in Denmark, followed by MASM courses in London. At RECKU I was immediately put into the EXEC maintenance job and quickly took over doing the gens, which were done nearly daily, constantly applying fixes, either PCRs received via telex or self-developed fixes. We also all had to analyze the many dumps and I was soon after introduced to dump analysis by Claus, my primary teacher. Claus told me to look out for two things that so often were done wrong then: 1) The JGD instruction address field requires that you specifically code it with User or EXEC register in EXEC code and 2) Third-word loads give sign-extension! I learned the trade of analyzing dumps at RECKU – the hard way, by looking at many, many dumps!

On top of this we had numerous onsite support people from abroad, for longer and shorter periods. These people were part of “the flying circus”, managed out of London by Jaques Arch. Some of these people became particularly good friends [we spent many hours together], especially Ian Hepburn and Jim Babcock.

The system had numerous problems and stops. The worst week we had 46 different stops – either hardware or software. The Centre had a centrally placed lighting system showing the system status: green meant up and available, yellow meant going down or in recovery, and red meant down. In that way the students could tell if they should stay around and wait, or they might as well come back later. When the system was down there was a few standard reasons for being down: CPU, IO, memory, software, etc. In bad periods they also had a sign with the reason why the system was up: Pure luck!



In the beginning, all student tasks were batched with a remote card reader outside the machine room and then the resulting print delivered the same place later. The card reader and normal printer was a 9300, controlled through a special program called SITE-SP (Slave Interface To Exec – System Program). Only the RECKU staff and we had terminals, in 1974 four U100s that we had to share amongst us, placed in a terminal room, not at our desks. Later, more terminals were added, and the number of demand-runs grew.

Why do I know it was one of the busiest systems? Because we were the first to hit a problem, which was caused by having many demand terminals active. The code in those days used to mark the address of the demand terminal control table [in EXPOOL] as negative when the demand session was in termination. But then one day our EXPOOL grew so much that the control table address exceeded 0400,000 – or a negative number when working with an 18-bit field.

The RECKU system had some excellent mass storage peripherals of the day, primarily three FH432s and two FH1782 drums, and then a FASTRAND and 8414 disks for “large scale” mass storage. The FH (Flying Head) drums were amazing devices with the FH432 having an access time of 4.3ms [but only room for 262K words – or 1MB.] The FH drums had vertical drums running at 7200 RPM and the read/write heads were flying on the air from the rotation (hence the name); it took many years before we had disk subsystems that could outperform these devices, and then only through caching. The tapes were Uni-servo 9C with a smart little unit between the controllers and the units allowing you to patch any unit as any device number, so you sometimes did not have to dismount tapes.

When I started, we were running EXEC 31.126 and soon after upgraded to 31.159C, as noted giving numerous problems and stops. The great breakthrough came with 31.244E, which should turn out to be the start of the very stable EXECs that we have since learned to love and appreciate – and take for granted.

The EXEC 31.244E master tape was hand-carried directly from Roseville to RECKU by Jim Babcock. We generated the EXEC, booted it, and it worked surprisingly well. We were enormously proud and the RECKU operations manager, a very demanding person called PRA [I believe his real name was Peter Rasmussen] could not believe it. But, after a few days PRA called us as the system was hanging [no stop, but no response to console inputs].

Somehow Jim managed to identify the problem and generated a patch for this. Since the console was hung, we could not apply the patch with the usual Console-key-ins, but Jim knew that we could apply patches through the card reader and prepared the patch cards. He walked into the machine room and read them in; it did not work immediately, and the situation was tense, but Jim quietly said to wait for six seconds – some very, very long seconds. Then the system suddenly took off again and Jim could leave the machine room in triumph! The EXEC 31.244E stayed up for nearly a week, unheard of in those days!

Jim Babcock stayed at RECKU for quite a while.

I only stayed on site at RECKU for 11 months, but there were numerous memories from that time. One of the more amusing was a trick we played with one of the RECKU system people. In the beginning, we had old U100s, which were equipped with rubber keys, which required that you hit the keys rather hard. Later they were replaced by terminals with spring-supported keys which only required a light touch. One person kept hitting the keys awfully hard and that was rather noisy and bothered other terminal users quite a bit. So, we decided to “teach him a lesson” by creating a special version of the ED editor that would recognize his user-id and then for each input came back and said “ouch”. He got upset – but he learned to use a softer touch.



Another special event was the day FASTRAND was taken out. FASTRAND is the unit that has defined the file addressing structure (sector, track, position) in the OS2200. It was a giant device of 2½ tons with two massive, silvery drums, lying above each other. It was installed in the basement, and it was impossible to take it out through the building [I do not know how it came in], so they had to make a big hole in the building wall. The FASTRAND was slowly lifted from its base, supported by wooden pallets until it reached the level of the ground outside. Here a giant forklift took the FASTRAND and very slowly drove it backwards outside so it could be lifted onto a flat-bed-wagon. It was hard work – and it even happened to be one of the hottest days that summer.

I came back to RECKU (or UNI-C) a couple of times when larger tasks had to be performed. One such event was the 1976 installation of the 2<sup>nd</sup> 1100/80 in Europe – the old 1100/82, max 2x2, with the white maintenance panel. From that installation two special events come to my mind. At first, we were not able to get the system booted, as it stalled when the first console message was issued; it turned out to be a faulty microswitch in the console page-writer [a printer copying all console traffic, a critical and required unit in those days]. After long analysis we found that even though the page-writer had correctly printed the message and returned the print wheel for the next line, but the EXEC was not notified due to a missing interrupt due to the faulty switch; a tiny, simple component – massive impact! When the system was then running and made available to the RECKU system programmers, Steen Jarbøel decided he wanted to test out the new instruction set, such as the TRA (Test Relative Address); but he did not know how to set up the registers and executed the instruction with the registers set to zero. This had the effect that the CPU just stopped – and then was kicked again on the next interrupt, e. g. from I/O, but in a rather undefined state causing a strange mixture of stops. We found out what happened by looking at the maintenance panel, where you could see the CPU going red for a short while before being started again – and combined with Steen admitting that it seems to happen whenever he ran his new test program.

The system was later upgraded to 1100/92, as far as I remember without special incidents. In 1988 I was also involved when a new system was set up [now in the computer rooms at UniC in Lyngby], a 2200/444 – a very large 2200/400. This was a different machine from the old 1100s and quite a few RECKU's special programs and handlers required some upgrades.

## BURMEISTER & WAIN (B&W)

In November 1974, I was transferred to B&W, a major shipyard and at that time the biggest employer in Denmark. The company also made the best and biggest diesel engines in the world and owned numerous patents for ship [and other large scale] engines. The company was in numerous locations in Copenhagen, but mainly on Refshaleøen (shipyard) and Christianshavn [diesel engine factory and administration].

We, Unisys and B&W support staff, had an office in the corner of the ground floor of the main administration building [for some reason called *Ørkenfortet*, *Desert Fortress*] with the machine room right behind and some equipment in the basement. We used the B&W facilities, such as canteen, but we had to adhere to the “code of conduct” there, which meant that there were certain tables where you should not sit, as they were “owned” by special groups.

The Unisys staff at B&W was more mixed than at RECKU and most people came onsite only when requested, such as John Iversen [and later Niels Gebauer] for COBOL issues. It was also here I first met my other mentor: Tore Ask. Tore was a Norwegian and a very experienced OS2200 person, having been involved with all the large Norwegian sites at the time. We also had consultants onsite for more application-oriented support, such as Arve Moss, who was a specialist in UNIS. B&W also had their own staff, primarily Finn Rasmussen [later SAS Data.]



I was sent to B&W because they had to upgrade their 1106I systems to an 1106II (going from 131K single access memory to 2x65K dual access interleaved and faster memories!). They also had an 1107 system running EXEC2 – and the 1106 also ran EXEC2 from time to time. For batch reader and printer, they used a 1004 system, a physically big (but small capacity) computer, programmed with some large plug-boards. It was also used for some offline computing, one being to read some paper-tape payroll information using a RC2000 [from Danish manufacturer, Regnecentralen – the only one that could read the square-holed paper tapes used] onto a tape, using a Potter tape unit. This unit had a weakness in that it often failed to generate an interrupt when the tape reached load point. But one day someone happened to bang his pipe against the unit [yes, you could smoke in machine rooms at that time] – and that made it give the load point interrupt; so, from that day the pipe was lying on top of the unit and a cross marked where you should bang the pipe.

Before the 1106II was brought in we had to take out the 1107 and its peripherals (the boot device for EXEC2 was a FH880). It was quite a slow task, one reason being because the tape units, Uniservo 2A (or 3A), were equipped with gold-plated connectors, and that had to be preserved and removed manually.

The move from the 1106I to the 1106II had one major issue. In those days you did not have room for [power or cooling or whatever] multiple systems and we probably also had to switch disk equipment. So, a SAVE ALL [using SECURE in those days] to tape was done, the 1106I and its equipment removed, the 1106II and its equipment installed and then the system should be loaded. However, the SECURE LOAD ALL [or REGISTER DIRECTORY] failed! That was when Tore showed his mastery and sat down and analyzed the huge, printed dump of SECURE – and found the problem; Tore also knew that SECURE had a built-in source level patch capability, so with added patch cards he managed to get the system loaded and up and running. True magic!

Another part of OS2200 that I got deeply involved with at B&W was Checkpoint/Restart. B&W had some complicated mathematical calculations [I believe to calculate stress in diesel engines] and they estimated that they would run for weeks – a lot more than the average MTBS (mean time between stops) at that time. Since these calculations were iterations, it was possible for a trained mathematician to help the process. Therefore, we set up the task to run for 8 hours, typically overnight, and then take a checkpoint - at that point written to tape. A mathematician would then look at the intermediate results and change the calculated data in the right direction to help the process. When restarted we would then have gained several hours, using careful re-assignment of files to the run [some from system, some from checkpoint – it took several attempts to get that right]. This turned out to reveal a few shortcomings in our EXEC. I remember two specific cases: 1) when the number of preassigned checkpoint tapes were exhausted the pointer to the next reel was zero and the EXEC looped trying to calculate a buffer address for the next reel and 2) when the calculations were finally done, and the results should be sent to a print file it failed as the restart process had assigned the print file with the default size (128 TRKs those days) instead of the normal print file size of 3000 TRKs.

B&W also had some remote equipment (DCT2000 remote batch reader/printer) out by the shipyard, so we had to go there from time to time – and had the opportunity to watch the giant shipyard a little from inside. It was through these visits that I learned about the shipyard and could talk about it at Diesel House, the reason for writing this story. But I missed seeing something that I was told to be a unique event: when they put together two parts of the giant cylinders of the diesel engines. It should have been quite a sight to stand inside the cylinder afterwards and not being able to feel the point where the parts were put together – so precise were the skilled staff of the engine factory.

## F. L. SMIDTH

Just before I joined Unisys, they had won a great contract to supply the main IT system for F. L. Smidth, a huge, famous engineering company, selling large factories all over the world, especially to produce concrete. In this contract, we should replace a CDC 3500 system, so the first part of the project was to convert all CDC programs and databases to OS2200 COBOL, DMS and TIP – and this was my first encounter with a transaction environment. The conversion project was led by Egan Larsen, but in the summer of 1975, I was put in charge of the software support as the focus was now more on operational aspects, performance and system problem analysis and solution.

Throughout the years many Unisys people were also associated with the FLS project, for shorter or longer periods. First to mind come our two “famous ladies” Bodil Skovgaard and Gerd Jacoby. In a world dominated by men these two wonderful ladies stood out and wherever you travelled in those days you could be sure that of all the people working in Unisys Denmark then they were the ones known all over. Bodil was the expert on TIP, with a background at Air Canada and SAS, whereas Gerd Jacoby, who had also been at SAS, was the DMS database expert.

The FLS system was an 1108 and it was installed in a new purpose-built computer building on Ramsingsvej in Valby, so compared to RECKU and B&W, where the computers were installed in existing older buildings, this was a genuinely nice facility. But that did not necessarily make them perfect. Throughout the years the FLS system had several problems and during one analysis it was revealed that the grounding system was constructed such that some false electrical capacities could be introduced; computers were overly sensitive beasts in those days. Due to unexplainable problems we also once had to rush a backup 1108 from Germany to FLS and install it in the back of the room, but it was never used.

FLS is a huge industrial conglomerate and around the main office in Valby were several company units, such as Valby Maskinfabrik, which produced the most special units needed in a FLS delivery. Whenever FLS had to build a concrete factory somewhere, it was often placed where the required raw material was available, so they often ended up in a very deserted and remote area. To build a factory, there FLS often had to first build a harbor for shipping the large parts and then a road that was big enough to carry the heavy loads. These were huge projects.

Among the central building was a tall office tower with a canteen on the ground floor and that is where we had our lunches. Apart from being ridiculously cheap [subsidized] the most remarkable was that there was a tap – with free beer! The canteen also had its special characters, such as a man that was nick-named “the big grey elevator-looker” by Gerd and Bodil. They gave him that name because whenever [they felt] they took the elevator he was there, and contrary to most people he was always staring at them. But he was also a frequent guest in the canteen, always there when we were there [early or late] and he always went up and took one sandwich, went to his table, and ate it, then back again for one more sandwich, possibly sitting himself now at a different table, and so on. We always wondered what his real job was.

When I came to FLS we had to pass the acceptance tests, such as stability and performance. One of the special requirements was that we had to prove that the system could run with batch, demand, and TIP, including DMS, in 131K words (524KB). The system had 262K words, but this was somehow a contractual item. Nevertheless, we did it – for just a truly short period. When you think about the amount of memory on systems today this seems completely unthinkable.

Over the years we hit numerous problems and due to FLS (and Ole Mørk) demands they often became crises. In the early days, all terminal handling was done within the EXEC (had its own Uniscope handler) and later this was moved into CMS – and I remember numerous situations of problems in these areas.





We generally fixed the problems ourselves, but sometimes we had to call in expert, sometimes only to show FLS some action. In one of those situations, we played a little trick and told FLS that we had called in an expert from the UK. In fact, we just went out to the SAS account (close to the airport) and picked up an English-speaking colleague, Peter Firth who was a good CMS man, and brought him to FLS – just to buy us some time and peace to work on the problem. In those days, the systems also had numerous hardware problems and I spent many hours troubleshooting with the CEs, mostly Jørgen Nielsen. We often needed outside help, some more special than others. One guy, flown in from Eastern Europe, fixed the problem!

One of the biggest challenges at FLS was that we were one of the first sites to move up to the EXEC-integrated TIP [in EXEC 33]. Before that time TIP was a separate package and all TIP files were allocated on specific removable disk packs, completely owned, and controlled by TIP. With EXEC 33 TIP files were mapped into EXEC files and co-existed with all other files in the system. This was a major challenge and took many long days [and nights] to get it working properly. In this I worked closely with Bodil Skovgaard.

We also had a few problems because we were one of the first sites to seriously use tape labelling – and then learned that you often see problems within areas which are not in daily use in the development Centre in Roseville. One rule at FLS was that the operator was not allowed to mount a write ring [a physical plastic ring mounted on the back of the reel, sensed by a switch to indicate that the tape was writeable] in a tape. The available output tapes were brought into the computer room by tape administrators and called by programs labelled blank tapes.

What happened from time to time was that one did not have the write-ring mounted [or could not be detected] and when the tape was opened the tape label was read and then should be rewritten with file information – but that failed due to the missing ring. Now the reel was replaced by the next reel in the stack and the I/O completed – but now writing the label-ID of the first reel onto the second reel. Several times we had to read many, many tapes to find the one with the mismatch between internal and external label.

Another problem area was the instability of disks in those days. FLS for most of the time used 8440 disks (two spindles above each other, behind glass windows). They had their problems, but what annoyed FLS was that the problems seemed to surface at the worst possible times. To try to catch problems as early as possible they asked me to write one of the most complicated system programs I have ever written. The HEADEx program ran every fifteen minutes and by using multiple activities and careful disk absolute addressing it made sure that every disk head on every disk surface was tested and continued with other heads no matter what problem it may encounter, reporting any issue as soon as they occurred.

Due to numerous MFD problems, I developed my MFDSCH program while at FLS; a program to analyze the MFD without relying on a correct structure and pointers. It came to incredibly good use during a serious problem with the 8440 disks, requiring some special action. During a boot, the system came up and reported a bad MFD link and could not recover all files, of course just before the file backup should run. From the reports from MFDSCH [run regularly], I could tell what was wrong with the MFD link, so FLS asked me to try to fix it. By mounting the bad [fixed] disk as removable disk on a test system, by using Absolute read and write [possible in those days – when checks were patched in the EXEC] and by using the head/cylinder address displayed on the disk CU to hit the right spot I managed to repair the link and when remounted the system was recovered. I earned a lot of “points” that night – and the respect of Garth McClellan, who helped me [by phone] during the process, but never thought it could be done!

After some years FLS needed a bigger system and after long discussions and benchmarks (see more about these later) they selected a 1100/42 system with large memories (of the time), with peripheral units shared with the 1108 using an SPU. This was a very unusual configuration but worked well. As this was also the time the spare 1108 was onsite, we then had one of the largest concentrations of 1100 computers in one room in Europe.

FLS was a particularly good assignment for me. I learned to handle the full suite of Unisys 2200 products in use [except DMS, still a hole in my knowledge] and built up a large contact network in many areas of our software [and some hardware too.] I also learned to handle a rather demanding customer.

FLS was also the first site where I got involved in software qualifications. This started with the move to EXEC 33 with its integrated TIP, but over the years I got involved with tests of many products – as OS2200 products in those days were released individually and not as the “integrated stack” [with many different names, such as SB, HMP, CP, etc.] that we have today – and rightly is proud of it. This helped me increase my personal network of people, both all over Europe, but also some people in Roseville. In those days we did not have a database of fixes/PCRs that we could investigate. Instead, we had to call the support Centre in London, where Les Shelley, as the only one, had access to the PCR database in Roseville; he would then send us the corrections, either on a telex or a tape, for a larger number of fixes.

Having to call Les meant that he had an incredibly good feeling for all problems encountered in Europe, so very often he could reply with his “I have a fix for that!”. Les was a great guy.

I worked initially at FLS for nearly five years – and a lot more hours than the normal working hours of five years! Since this was my first long term assignment, I also here had the first feeling of being closer to the customer than Unisys, both workwise and on the personal level, sometimes resulting in a conflict of interest, where you sometimes saw issues more from the customer point of view. Jørgen Sindahl (Jensen) was also very sharp in databases and their DMS expert, and when you asked Jørgen a technical question you always got a very, very full, and precise description [always starting from “Adam and Eve”.] Kailas was from Pakistan and a small man; he took care of the basic software products and was my closest counterpart. Kailas was overly cautious, which some people considered slow, but Kailas was just detailed, as you had to be in the OS2200 world.

I left FLS in the summer of 1980 to start my assignment in London. This was also an end of an era in Unisys 2200 world, as this was the beginning of the “cheap” systems like 1100/60, where the cost of the system could not carry the cost of a “free” Unisys System Analyst and CEs with every system. So, it was probably good that I switched to another job, as that had been difficult to convince FLS about.

I later worked on FLS projects a few more times. After I moved, they had a few additional benchmarks and for all of these I was called in due to my deep knowledge of the systems. I also had a brief assignment in 1997 when FLS was migrating to their last OS2200 system, an IX4400.

Before leaving the FLS part of the story, I should mention that my university project program LACAP was taken to real and serious use at FLS. Someone at *Petershaab Maskinfabrik* [owned by FLS] ran into the usual problem of exceeding the capacity of the ANP3 program. Therefore, I implemented my LACAP program on the FLS Unisys 1108 system, and it solved their problem. The first time I ran it I forgot all about the many JCL statements required to run it on the IBM system – but it worked anyway as Unisys FORTRAN handled all this internally!



## OTHER DANISH SITES

There have been numerous other OS2200 sites in Denmark, and I have been involved with many of them on few or more occasions. I have not been onsite for long periods at the sites below, but small visits can also bear some memories, good or bad. Many of these belong chronologically later in my story, but I cover them all here for practical reasons.

In the early 1980s Unisys had great success in selling numerous 1100/60 systems in Denmark, some as pure OS2200 systems, but quite a few of them sold because the 1100s were the system required to run MAPPER. As I was living in England in those years I was not deeply involved with these sites, but I was asked to come home during the summer of 1981 as many systems had to be installed in a short period.

### Dansk Data-service

**Dansk Dataservice (DD)** was quite a big site in the 1970s, a service bureau, the “cloud installations” of those days, providing IT services for a broad range of customers who could not or would not have their own IT department. I did not have a lot to do with DD, but was onsite several times when they needed help with the EXEC.

### Danish State Railways (DSB Bundesfahr)

The **Danish State Railways (DSB)** was a particularly important and critical customer for Unisys Denmark. It was also technically advanced in several areas, so we had a larger group of people there: Palle Hansen, Bo Sørensen, Thomas Dynesen, Jan E Jensen, Michael Deneyer, Jesper Ingvardtsen, and Per Bo Rasmussen, plus others from time to time. They could handle nearly everything, but a few times I was called to DSB to look at special problems or just have technical discussions.

DSB was primarily a COBOL-UDS/DMS/RDMS user, early adopter of OLTP, early implementer of Extended Mode – and maybe the only site to seriously use the OSI-suite of products. Being a state-owned company, they were sometimes forced to use “official” standards (like OSI) instead of de facto standards. All of this gave a few challenges and kept our people there busy – and made them very good OS2200 people; fortunately, we still have some of them around.

### DataCentre ved Odense Universitet

“**DataCentre ved Odense Universitet**” (DOU), the regional computer Centre at the University of Odense, was one of the bigger sites in Denmark at the time. It was also the breeding ground for one of the key OS2200 people in Denmark, Jan E. Jensen. The primary person on-site was Jan Kennet Larsen, who also worked as an assistant professor at the University and handled most things; but I was involved with several initial activities there.

One other remarkable incident at DOU was a problem with the cooling system or a false alarm, where they ended up sprinkling water over the machine – with salt added to the water; not a good cocktail for electronics. It took the CEs quite a while to clean the system and bring it back up again.

### Synoptik

Synoptik is a chain of stores selling spectacles and lenses, with stores all over Denmark, in Norway and in Germany. They were a very heavy MAPPER user, which required the use of an XPC (in caching mode) to deliver sufficient performance. Synoptik was also a rather complicated customer, primarily because of their head IT person Jan Andersen. One day they hit a problem with the XPC interface resulting in the loss of all cached data. As it turned out that Synoptik did not have proper operational procedures and no backups it was quite a challenge to recover the system and it required the work of lots of Unisys people to get them going again.

My only other encounter with Synoptik was a performance study that I carried out together with Gary Tinker [then in Unisys UK.] I believe we produced a very professional and correct report, but Jan Andersen did not like the conclusions and recommendations – so from that point I was persona-non-grata at Synoptik.

Synoptik was, as far as we know, one of the last OS2200 users in Denmark. When support of their IX4800 ended, they hired an ex-Unisys CE as their technical consultant, stopped paying maintenance fees to Unisys on their one-time-charge licensees and had no official support. But, of course, the system kept running for quite a while. We also heard that at some point they bought a second-hand newer system, without understanding that there was no way that their old EXEC would ever run on that system.

### Meteorologisk Institut

For a few years we also ran an OS2200 system at the Meteorological Institute of Denmark, a 1100/80 system – with an attached vector processor [a very rare beast indeed.] I was only there briefly during the installation, but several good Unisys OS2200 people got involved with this system. This is also where Mikael Rattenborg got involved with OS2200 systems.

Mikael is the sharpest person I ever worked with! He had a degree in Meteorology, with the highest possible marks. but got interested in our systems. When the 1100/80 was shut down he came to Unisys and worked with me at SAS.

## **BENCHMARKS**

Benchmarks were a common way of selecting a new IT system in the 1970s. In the purest way, a Benchmark was a well-defined test run of a typical subset of the customer's workload and then the vendors would implement a solution to run that workload, present and run it. Then ideally, the one who ran it the fastest would win the contract. In practice, this would never hold, as many other factors could affect a decision; but I believe at some point the rule was, that if a customer asked for a Benchmark, which could be rather expensive to run, then they were obliged to choose based on the result.

Benchmarks were often quite complicated to set up and run. It would often be a project nearly the size of implementing the system; but if you made a good Benchmark, then that would often help a long way for the actual implementation later.

I was involved with several Benchmarks in my early years. They were challenging and difficult to solve with unbelievable work hours – but they were also fun and very rewarding to be part of.

### DOU Benchmark

My first Benchmark, in 1975/76, took place in our Benchmark Centre in London, on Westbourne Terrace, near Paddington Station. DOU (**D**ataCentret ved **O**dense **U**niversitet), the computer center at Odense University, needed a general-purpose system for students.

The proposed system was a 1100/22 and the test was a mixture of demand and batch. The main technical effort here was to develop the communication simulation drivers and scripts to simulate the many demand users the system should be able to handle. OS2200 used to have special software modules for this and a special simulation program CS1100 was used to feed the system and handle the responses, controlled by a script, called TCL. This was not commonly used and gave quite a few problems – but in the end, we developed some scripts that I used for many years to follow when testing OS2200 – until the simulation modules were removed.



One special requirement of the test was that the input set of cards [yes – in those days the batch runs still entered the system by reading in punch cards] should be read in over and over for the duration of the test. What the test designers had not realized was that our fantastic 0516 card reader could read 1000 cards per minute, meaning that the runs got into backlog at an unbelievable rate. When they saw this, they realized this was meaningless and OK-ed that we stopped reading after a while.

Another big challenge was that DOU had one key person who was a dedicated APL user [APL was supported in OS2200 then], so we had to develop simulation for that. This was complicated as APL uses all kinds of “strange” command, short-cuts, etc. that normal terminal communication never used.

We managed to get it all working – and won the contract. On the actual test night, we ran into numerous problems, and we had to stay at the Centre very late. But DOU still accepted the result and continued as a particularly good customer for Unisys Denmark for many years, as told above.

### Tjæreborg Benchmark

My next Benchmark (in 1977-78) was also the biggest I ever took part in – and it gave me a lot of fantastic experiences, good and bad. It was for Tjæreborg Rejser, at that point one of the biggest travel agencies in Denmark. The company itself was also a special story. It was owned and run by a priest (Eilif Krogager) out of the small village of Tjæreborg, located in Western Jutland, far away from Copenhagen and the other bigger cities in Denmark. It had started by the priest taking some local people on a bus trip to Southern Europe and then grew into a large company with their own charter airline, Sterling Airways.

Tjæreborg already had an IT system [from CDC], but they needed a bigger system as they were expanding. Long before the Benchmark there were numerous discussions and meetings and I still clearly remember a visit to Tjæreborg. I will never forget one remark, where we were talking about the special concerns that a system also had to cater for. They told us about the fact that people tend to live over ability on vacation trips, and consequently there was a higher-than-average death rate amongst charter passenger; so, one special concern was that some people had to be returned in a “single room with a lid” as they expressed it, i.e. a coffin.

The Benchmark itself was in my view very realistic [probably the best I have ever seen.] We were provided with a full database design and an Audit tape of a full day’s transactions. We then had to extract some specific records from the Audit tape and replay certain passenger transactions of that day. For the live Benchmark we would then be provided with an Audit tape for a comparable day and had to run our programs for that day; it turned out that they knew their data well as the two tapes really were comparable.

With the database definition, we then started a long preparation process, which was an experience. We were proposing our brand new 1100/82 system, and we were lucky enough that one had already been installed at SAS, at that time only being used for operator training and very initial USAS development; so, we were able to borrow it for our development. The software solution was based on TIP and DMS, also new and untested in those days. A talented group was set up and based out at SAS and we started developing the programs, set up runs, etc. I do not remember all the people of the group, but key members were Torben Brandt, sales; Tore Ask, software manager; Bodil Skovgaard Toftager, TIP; Gerd Jacoby, DMS; Niels Gebauer, ACOB programs; and me, EXEC, runs, and utilities). We also had ad-hoc help from people like Andy Øverby [from Norway], Arve Moss, Dan Ulmark (from Sweden) – and probably a lot more.



Established in 1980

We worked awfully long hours and there were lots of problems, discussions [at least one leading to screaming and tears] and conflicts with the SAS staff because we used the system to a level where we nearly took it over.

But in the end, we had it working and it was time to move the effort to a Benchmark Centre – and the only one having an 1100/80 was the main Centre in Eagan, Minnesota [just south of Minneapolis/St Paul.] Since I was responsible for the EXEC and thereby the basic set up, I was sent to Eagan a week before the rest of the team, in mid-January 1978.

That trip was something special – probably the most complicated I ever did. First, it was my first trip to U.S., so I was very excited. I should fly with SAS via Montreal to Chicago and then on to Minneapolis/St. Paul. But I hit problems as soon as I reached the airport as I was told that my flight was cancelled, but I was rebooked via New York instead. So far so good – and the first flight worked well. I landed in New York, and I had my first experience with U.S. Immigration Control. We walked through some long corridors [SAS then flew to JFK] and then we hit a queue, long before the Immigration area was in sight. It took forever to go through and after Immigration it became clear why. Everybody had to open all their suitcases for Customs inspection! I have never had that happen again, so I wonder what was special that day. I was carrying some computer tapes with all our programs and before leaving we had discussed whether I needed special customs paper for those tapes; I am glad I had those papers ready. The whole process took 4½ hours, so when I got through, I had missed my connection to Minneapolis; but SAS had rebooked me, and I was told to take a bus to the Northwest Airlines [probably Northwest Orient in those days] terminal.

It was horrible weather and the route between the terminals was a mess with cars and buses blocking each other all the way. The result was that the inter-terminal transfer took so long that I missed my new connection. I later learned that the Northwest terminal was only one or two buildings from the international terminal, so I could easily have walked the distance clockwise instead of taking the bus counterclockwise all the way past all the other terminals.

So, when I came to the Northwest counter, they had rebooked me once again – an hour or so later. I walked around but I did not dare sit down as I was very tired at this point. We then got onboard the plane – only to learn that we would be delayed due to some problem with the onboard intercom system. By then I fell asleep in the seat.

I woke up when we touched down for a short stop in Detroit. We stayed in the plane, but we had to wait quite a while because it was snowing very heavily, and they had to clear the runway between each take off! But eventually, at 3 AM local time [10AM Danish time – the next day], I reached Minneapolis. Fortunately, my bags were there [I had hand-carried them in the bus in JFK, so they had followed me] and I could walk up to Hertz [fortunately they were open] to get a car – again a first time for me. It was a helpful lady that explained the route in detail, realizing this was my first encounter with American roads. I also remember my question to her: how to start a car at -30°F – as it was that cold! If you have forgotten how carburetor cars worked the answer was: push down the accelerator once, wait 30 seconds and then start it; it worked! Fortunately, I easily found the Marriott hotel on I494 and Cedar [a short drive], got a room – and fell asleep immediately.

Next morning the challenge was to find my Benchmark coordinator [at that time each Benchmark was assigned a coordinator, a single person, who took care of all our needs – worked very well!] as all our meeting arrangements of course had fallen apart – remember we had no cell phones then! I walked down to the reception and found a message to call Stan Simmons, but he did not answer. By pure luck I found him in the telephone booth next to where I tried calling him – when he mentioned my name to someone! So now we were ready to go to the Benchmark Centre.



The Eagan Benchmark Centre was then in the impressive building at the corner of Yankee Doodle Road and Pilot Knob Road [still there, now Lockheed Martin.] As the building also contained the Federal Division security was very tight and I was immediately taken to a Security Officer, where I was told where to go and where not to go; that is when I first found myself labelled as “an alien”, since I was a non-U.S. citizen. Since the 1100/80 was the pride of the Centre it was on the first floor with glass windows to a central open area.

In the computer Centre, I met our assigned operator [ours, just like the coordinator – again a very good idea], Jenny Andersson [I believe]. The operator and I had spent a lot of time together. At first, she expected that it would require some effort to adjust our runs and programs to the 1100/80, but she was so surprised to hear we already had prepared on an 1100/80 as Eagan thought they had one of the first systems in the world. So, I had the correct EXEC, runs and programs right away. But when we wanted to load the generated version of TIPUTIL, we had serious problems. As soon as the tape was mounted the whole system was hanging. This took endless hours of dump analysis and testing, involving people from development in Roseville to sort it out. In the end it turned out to be a design flaw in the 5032 tape CU for the U3x tapes, when the distance between the load point marker and the P-burst ID was too long; we got another tape and a workaround to the EXEC. When the others arrived, I was ready.

The other members travelling to Eagan were Bodil and Tore with Torben coming just before the actual test run. We now spent many, many hours running the test, improving the programs and database layout and optimizing the file placement for best performance. The Benchmark itself, run on 9<sup>th</sup> February 1978, was an outstanding success. Everything worked – and we beat everybody else by running the Benchmark in 36:12 [I am not 100% sure – I have forgotten; I thought I would never forget how long it took]. Tjæreborg had visited Burroughs in Detroit the day before, and they had major problems in extracting the data from the real Audit tape. So as an extra test they wanted to time that extract too; we ran it in 32 seconds – it had taken hours at Burroughs [it really looks fantastic when you read a 7-track 556bpi tape and write it to a 9-track 6250bpi tape – one spins, the other hardly moves] – I wrote that program! Everything else worked fine – in fact, the real Benchmark and its disk performance balance were better than any test run. Tjæreborg had heard that we were doing very well – so they trusted us enough to have booked dinner with CDC the same day we ran the Benchmark!

Yes – we won that Benchmark. Both Burroughs and CDC [using two different systems] used around 50 to 55 minutes for the test, Data General quite a bit more and so the contract should be won! On the night of the final Tjæreborg internal discussions, they also called Torben around midnight and told him that we had won. But then at 3AM they called Torben again and said they had now chosen Data General. The next morning, they came to us in Copenhagen and tried to explain why they changed their mind, giving some strange explanation about better system architecture with the Data General systems, something we never understood or accepted. Years later Poul Spring met one of the Tjæreborg key people and got a more trustworthy explanation: after making the decision the Data General fan [Svend Ladegaard] brought up a question of where they should put the system in the transition period. The 1100/80 was large, tall cabinets and it would block the windows if they were installed in the only spare space available – whereas the Data General cabinets were low and would not block the light! Totally exhausted after weeks of evaluation and discussions I guess it is tempting to jump to an easy solution. I do not think the new system ever became operational; in the years that followed Tjæreborg merged with Spies, another travel agency – and then I lost track.

### FLS 1100/40-SPU Demo

In early 1979 we went back to Eagan for a “Benchmark” for my own site, FLS. This was not really a Benchmark as we had no competition; but to get a contract we had to prove to FLS that the proposed solution could work.

The task here was twofold. First, we had to prove that all the FLS programs, developed for the 1108, would run without modifications on the 1100/40, and secondly that the two systems [1108 and 1100/40] could share peripherals using the SPU (System Partition Unit) designed for 1100/40 to designate units to one system or the other. This was an early attempt to define a proper DR (Disaster Recovery) system. Here we had great use of the CS1100 communication simulation scripts developed for the DOU Benchmark as that allowed us to run a representative load. We had an extra task of developing new scripts to run MAPPER, which, like APL, has a rather different communication protocol. Finally, we had to demonstrate new DMS features and the use of QLP as well as RPS. The test worked well – and FLS signed the contract for the 1100/40 with SPU.

### EDB-skolerne Benchmark

This Benchmark was run in April 1980, also at the Benchmark Centre in Eagan. The competition was for supplying computers to four regional IT schools in Denmark. [Copenhagen, Aarhus, Aalborg, and Odense, where the latter was a renewal of the system at DOU]. The systems offered were 1100/60s.

The test was again a mixture of demand, using CTS, and batch. Instead of using the TCL system [see DOU] we this time used the CNE system, a terminal simulator running on an 1100/40 system, which meant the system overhead was considerably reduced. But it was a new challenge to make this work. On the actual test night, we ran into numerous problems, and we had to stay at the Centre very late. We won this important contract and for many years IT professionals in Denmark were trained on Unisys systems, something that meant that Unisys computers were well known in Denmark for the following years.

### FLS Cache Disk Demo

My next visit to Eagan Benchmark Centre was in 1982 and was also for FLS, this time to show the performance improvement they could gain by installing Cache disk. These were the Unisys-developed Cache disk, using the 5056 controllers, controlled through software in the OS2200 EXEC. I had then left Denmark, but FLS had asked that I took part in this test, due to my year-long assignment there and deep knowledge of their system [FLS never got an onsite System Analyst after I left]; also, I was living in Minnesota at that time.

We used the same scripts and tools as on the previous tests, supplemented with some very heavy batch database runs. This demo was a little different as customer people, Ole Mørk, Jørgen Sindahl and some financial manager, took part in the preparation together with Tore and me; for one reason as they were just as interested in getting the deal as we were!

We had quite a few technical challenges in this test and spent exceedingly long hours at the Benchmark Centre. We also had to fight for test hours as our tests took hours to set up and ran quite a long time. In the end the actual test happened one night, where we, were called to the Centre as we could get a spare slot here and now. So, the test was run with only Ole and me awake, while the rest of the team slept in some sofas and chairs, totally exhausted.

Once the technical test was well completed, we suddenly faced some discussions about contract and delivery dates. We had some long non-technical talks in Roseville, but in the end, we managed to get some special arrangement, primarily because I in the end, in pure frustration, called Jørgen Rasmussen, a Dane working in some financial group in Blue Bell. Then we got that contract!



## FLS 1100/73 Demo

My last Benchmark in Eagan was in June 1984 where we had to rerun very much the same tests as above on an 1100/73 system. This demonstration started with a management presentation in Roseville – where the amazing Sperry PC was also demoed! Those were the days! Otherwise, I have no special memories or stories from that Benchmark.

## Unisys AIDS and COMUS

As noted above the “cheap” 1100/60 systems, hitting the market around 1980, could not carry the cost of an SA (System Analyst) assigned to each site. To then automate or simplify some of the tasks that the onsite SA used to do Unisys decided it was time to develop a standard software package for this. Unfortunately, Unisys U.S. and Unisys Europe could not agree on this, so two parallel products were developed, with very different focus and goal.

In the U. S. some management group sat down and defined the need and wrote a SOR (Statement of Requirement). I remember that one of the demands was a maximum size and that it should run on all systems, including a dual-PSR system like an 1108 [I had the challenge to prove that later]. But no existing development group had spare people for this, so the job was given to two university-graduates from Florida, Bill Slobotski and John Pokladnik.

They started with the SOR and developed a huge tool called COMUS, written in PL/1 on the OS2200!!! Over the years they managed to be able to support the generation of such complicated products as ED and FURPUR – but it never really worked. Harold Royalty of a U.S. customer wrote a wonderful article in the USE [the U.S. User Group] magazine about “Dr. Strangegen, or How I Learned to Stop Worrying and Hate COMUS” – a rephrase of the subtitle of the film “Dr. Strangelove”. Bill and John were OK, but they had no SA experience and never knew site requirements.

In Europe, the Unisys management decided to put together “the best site SAs” to produce the tool to replace themselves. A small group started this and reviewed quite a few existing local [country or site] tools, and they also looked at a rather elaborate tool used in Japan, called SMART. In the end they designed a database for fixes (CHGs) with a search/selection method based on a German document-retrieval system called UNIDAS and a generation tool very much like the Japanese SMART. The product was called AIDS (Automate Installation and Debugging System) – at that point in time an ok name.

With the design in hand, they decided to build up a larger organization which would also be a central support group. In that way the group grew, and I was asked to join. My management in Denmark was reluctant and allowed me to leave maximum one year – in the end I was involved with this project for 3½ years – and I had a fantastic time.

The team consisted of an incredible set of very talented people. In the beginning, the manager was Jim Davis, later Mick Daly and then Keith Binks. At its maximum, the group consisted of [sorry, I know I miss some – we were more than these]:

### DBA (COMUS Database):

- Seto Assilian, Dave Silber, Hans Ulrik Hansen, Ralph Platten STI (Gen part, later BUILD):
- Ian Hepburn, Mel Peterson, Ted Lambton, Mike Gleen, Doug O’Neale, me SQA (of AIDS):
- Gary Moerke, Joe Nolan, Barry Johnson, Rick Spence, Tony Foll TAG (Telephone Assistance):
- Les Shelley, Jørn (John) Loberg, Pat Lowney
- Onsite support group (“flying circus” – never in the office):

- Paul Cameron, Jim Douglas, Merv Griffiths, Jim Signorelli, Alan Bennett and quite a few more, with Paul Fonstad as their coordinator.

The group was established in a new office near Paddington station on North Wharf Road – and yes, there was a wharf there, as this was the very end of Grand Union Canal and there were barges moored under the building next to the parking lot.

## MOVING TO ENGLAND

I came to London in early July 1980 – and the family should come a month later, when I had found a place to live. I was installed in a small family hotel and so was Mike Gleen. Mike was an American, from Albany, New York, but had worked in the 1100/60 test group in Roseville and met some of the others from the AIDS group.

My years in England was really split in three parts: the development and support of AIDS, the migration to COMUS [partly living in Minnesota, U.S.A.] and the post-COMUS period of general field support.

### AIDS

I should warn that these next sections have quite a few technical terms, which can be difficult to understand if you have no OS2200 knowledge.

The background for and the groups in the AIDS development team is described above. I was assigned to the STI (**S**ystem **T**est and **I**ntegration) group, headed by Ian Hepburn. I have known Ian for many years [see under RECKU] and found it quite easy to work in his group. The group was responsible for the product generation run stream and the adaption of the skeletons to the defined AIDS standard for handling merging and applying corrections. As part of this process, we had to take all corrections from the development PCR database and manipulate them to fit AIDS standard and store them in the EAD version of the database to convert the release tapes to AIDS format release tapes.

The latter task was initially my job – and that was quite critical as this step had to be done after the official release and before the products would be available for the European customers. In converting the PCRs and re-applying them as CHGs we had to recreate each generation step done in the development Centre and then at the end compare our final product with the released version. For this process to work we were highly dependent on each PCR being correctly marked for internal relationship and system generation numbering, and quite a few times we had to go back to the development groups and had them explain their process and update the PCRs or whatever was required for the CHGs to merge correctly. With the time difference to the U.S. [and before the time of emails] this could delay the process. So, at release time [in those days products were released individually] I had some tense, long days.

Where the COMUS people in U.S. started with small products we started by handling the EXEC, the biggest. We did not change the normal EXEC generation process, but merely wrapped it in processes that would handle the merging of CHGs such that the final TCF could be passed to the EXEC generation process. EXEC was selected because that was where most sites needed to apply fixes (PCRs) or had local corrections in those days, so for AIDS to be a help it had to handle EXEC such that corrections could be maintained in the customer AIDS database. Products were gradually added, focusing on products that either had local code or where fixes most frequently had to be added.

One of the initial features of AIDS was that it would automatically detect what PCRs were applied from the local database and exclude them if they were contained in a new release.





The actual AIDS generation run stream was simple with an input set up part, a definition of several parameters defining the generation process and then executing the SMART program [name inherited from Japan]. This gave little opportunity to manipulate the process beyond what the parameters allowed. When we later migrated to “COMUS 2” this was changed, so that the run stream became more open and accessible. The initial AIDS generation run had to be manually edited and at some point, it was decided to make a program for setting up the run stream – the BLD processor [later converted to COMUS BUILD command].

BLD [and BUILD] is my baby, the legacy I leave behind in Unisys; and it is still in full use nearly 40 years later. BLD is my design and to large degree my development. As usual for a tool it started its life as an internal utility, because we probably did more AIDS generation than the rest of Europe put together. So, we used an available tool for it, namely CTS, which had a very powerful script language. When turning our procedures into an AIDS product we ran into the problem that CTS was not a free product (like AIDS had to be). We got around this problem by some very tricky changes to the CTS processor, so when called as BLD it could not allow normal CTS processing. Many “dirty tricks” were used to hide CTS and lock the user into the BLD environment. One was that all procedures had to have a version of /BLD, even though they were never called that way. Also, to prevent BLD revealing itself at CTS, if for some reason the BLD procedures lost control and would return to the standard CTS prompt (>->) then the simple prompt was replaced by a cryptic message, as far as I remember: “ERROR – Nesting loop detected in BLD”. Many of these tricky CTS changes were the brainchild of Doug O’Neale; and I must admit that we sometimes had great fun in finding “strange ways” of hiding CTS. Doug was a fantastic programmer and system analyst, but he is no longer with us as he suddenly died from an instant heart failure, in central Milton Keynes, only 48 years old.

In 1981 Unisys management finally decided that it made no sense of having two different products and after some internal fighting it was decided to merge the two products by calling the product COMUS (2R1), but the contents were AIDS. Name-wise this was a good choice, because this was just at the time that the horrible disease AIDS was discovered/identified in Africa. So as Joe Nolan cited it a few years later:

- In 1980: COMUS was a software product and AIDS was a software product.
- In 1980: It was better to have AIDS than COMUS.
- In 1987 COMUS is still a software product and AIDS has become a disease.
- In 1987 CLDC can at least claim that it is better to have COMUS than AIDS....

A project was then started to merge the two products. In the meantime, COMUS development was moved to Clear Lake, Iowa (CLDC in the quote above stands for **C**lear **L**ake **D**evelopment **C**enter). Unisys had taken over the facilities of a closed-down CB-radio factory [probably with huge subsidies to save jobs in this small town] and a few software product developments were moved from Roseville to Clear Lake, together with hardware final assembly and testing [the building was built for and ideal for static sensitive work]. The operation in Clear Lake was headed by Bob Krieger. When Unisys many years later decided to pull out again, the people in Clear Lake created their own company and took over some of the products that Unisys had pushed down there, primarily the performance analysis products. In that way Team Quest was born!

Several AIDS group members were sent to Clear Lake for shorter and longer periods, me included. We also took advantage of this “new product”, which would require some migration work both for COMUS and AIDS users anyway, to fix up some parts of AIDS that time has proven not to be too good. This was a great advantage for the product to have been extensively field tested for years, and maybe a reason why the product is still around.

Two of these improvements were the more open run stream [as previously mentioned] and that the BLD processor became integrated as the BUILD command. In this way we also got over the CTS issue as it was now locked inside. But it also gave us huge a challenge to MAP COMUS [also needed all ACOB inside] as the product had to be independent – and still had its old max size requirement. It was a hard task, but in the end Peter Bergh in Roseville managed to do it!

### COMUS

By the time, the updated common product programs and skeletons, mostly done by Doug O’Neale, were working it was time to address the BUILD command. For that I was sent to Clear Lake for the month of May 1982.

Clear Lake had 7,500 people mostly in the summer; in the winter many people, often retire, moved to Florida and as the name implies built around a large lake in Northern Iowa, near Mason City. Unisys was the largest employer in town.

The next part for me for COMUS was to adapt all the BLD procedures to BUILD subroutines, for all the 28 products at that point supported by AIDS or COMUS 1. This was a major task and should be done in cooperation with the development groups. This part of the project was run by Joe Nolan and me, with Ralph Platten helping on the UDS products. As this was a lengthy project Joe and I with our families moved to Minneapolis/St. Paul in July 1982.

The COMUS work was very interesting – and challenging. Our little “office” was set up in a secretary “cut-out” along the wall in Roseville building 2, instead of having two standard cubicles, which meant that Joe and I could share everything. Our small office and the project were nicknamed the “COMUS Factory”. We were assigned to the Release Management organization and reported to Willis Unke. But our task meant that we had to work together with all the development groups that had products released under AIDS and COMUS, and that was a large amount of people, giving me a kick-start to my rather large network within Unisys.

It was decided to take one product and make a model for that product to be replicated as much as possible for all other products. I do not know who picked ACOB, but it turned out to be quite a challenge. I remember from my early days that the people who could generate ACOB were very few and far between, in Denmark it was Niels Gebauer. We used the first time then to set up routines with as many options as possible that would satisfy the needs of the ACOB generation [ACOB was rather special in that it generated several variations of the compiler – and one of them had to have 57 errors during the MAP to be correct!] From the initial generation we started to build a library of common routines with comparable structure, mostly to make the BUILD process as identical as possible across all products, but also to ease the work of implementing 28 products.

The next product we addressed was the EXEC and we worked closely with the EXEC Utility group, then headed by Bill Gilbertson. Our implementation was not received equally well in all groups and accepted right away; nobody likes to have some “outside” solution dumped on them to replace their established procedures – and in this case delivered by “aliens” [non-U.S., non-Roseville people]. Several development groups clearly just put our solution on a shelf, but when release management soon after insisted that the products had to be released with COMUS BUILD routines then many groups took the quick solution and used our solution – and in that way many product generations are still very alike to this day.

The only products where we needed help were the UDS-suite of products; since Ralph Platten had done these for AIDS he came over and joined the group for a while for this implementation.



The only non-Roseville product we had to do was TELCON, so I travelled to Salt Lake City to work with the developers out there to set up the routines. It was such a pity that I had to leave Salt Lake City on a Friday afternoon as we had some critical work in the weekend – as this visit was in early December and the best powder snow was lying just outside the hotel window; what a waste not to try to ski one of the great resorts there!

One additional task I had in Roseville was to run COMUS on a 2-BDR machine, the 1108 S1 system in the back of the Roseville CSF computer Centre. Of course, it worked! So, I used an 1108 in 1982, the last time I ever ran on that magnificent machine, the one that the whole OS2200 product line live on! A big hurrah for the people who built the 1108 – and with it the “EXEC 8”.

#### EAD SUPPORT

When COMUS was released, and the initial problems were sorted out I was assigned to the general EAD field support team. I did not have many trips to customer sites, and spent most of the time in the Stonebridge Park office answering field questions, analyze reported problems, etc.

Here I need to say a few words about Stonebridge Park. At some point our “private office” near Paddington became too costly to run and we were moved to the bigger Unisys office at Stonebridge Park, where most of Unisys EAD and Unisys UK were based. It was two large bent buildings – in the middle of nowhere.

Most of my work in this period was follow-up and support for COMUS, and I only had a few on-site trips, which you can read about below. I therefore spent a lot of time testing out Integrated Recovery products and procedures, and this gave me two rather special support trips. One was a trip to South Africa [see below] and the other was a trip to Harrisburg, Pennsylvania [next to Three Mile Island]. I was in Roseville for some reason, but because of my IRV2 testing some manager asked if I could join a team on a trip to Masland Carpets, together with numerous Roseville people [amongst them the only person I have met that worked on a 1103]. We had to help Masland Carpets set up what was probably the first environment running the full suite of IRV2 products.

I owe my AIDS and COMUS period for the large network of Unisys people I now have, a very great contributor to the success I have felt in my professional life. A very warm thanks to all of you for all the help over the years; without my network I never would have made it!

One consequence of this was that I was invited to the “LDC reunions”. LDC was the London Development Centre, based at Westbourne Terrace. Several development projects have been run here, most famous the RTOS operating system, but in all cases I know of the product developments have always been moved back to the U.S. at some point. However, LDC was also the European benchmark Centre [as described earlier] and the central European computer Centre. We also used the LDC systems for all AIDS development, testing and the integration work. Because of this I spent a lot of time at LDC and knew most people there. So even though I was never part of any LDC development group I was invited to the “LDC reunion” of over 100 people on 11th October 2003. This was a great event and people came from near and far, quite a lot of people from the U.S., and Garth McClellan came all the way from Sydney, just for this. I was also invited to the repeat event in 2009, but unfortunately could not go.

## SCANDINAVIAN AIRLINE SYSTEM

My assignment at SAS (Scandinavian Airlines) was my longest and in the Unisys OS2200 world I will probably always be associated with SAS. I started at SAS when I returned to Denmark in December 1983 and was onsite for 18 years, until 2001 – with a few breaks – and with a few returns for short-term tasks a couple of times after that.

The tragic side of the story is that I was initially writing this part of my story on the day that I received final confirmation that the SAS Unisys systems would be finally shut down on 31st May 2013; ending a business relationship that lasted nearly 50 years.

SAS began their Unisys era with a 418 system in 1964. This was later replaced by a 494 system and grew to multiple 494s, running one of the most advanced airline systems in the world. With the termination of the 400-series SAS decided to go with Unisys OS2200 series and got an 1100/80 system in 1977. At the same time, they contracted for the Unisys not-yet fully-developed airline software package USAS and SAS played a major role in the final development of this product – and became the first airline, outside Japan, to go online with USAS on 8th March 1980. Alongside this SAS also developed their huge communication network based on Unisys equipment [3760s and then DCPs, set up in an advanced dual-system configuration].

This also was a front-runner in development of the TELCON communication software; but that was all set up before my time. Alongside these major systems SAS also over the years ran some smaller systems, such a V77-800 and U6000 Unix systems.

The Unisys organization at SAS was large and for many years acted as a subsidiary/country. In 1984 we had something like six managers, two secretaries, one financial person, six OS2200/USAS consultants/programmers [e. g. Karen and Garth McClellan, John Read], six TELCON programmers [e. g. Poul Tønder, Claus Muff Christensen] and a lot of hardware engineers, both for the 1100s and for the DCPs [e. g. Norman Errebo, Villy Laustsen, and many others]. I took over the job from Martin Løfgreen and Dennis Konrad, who passed on to me a very well-run site. But the common figures throughout all my years at SAS were Per Voss, who was responsible for sales, and Bent Strøh, who was there for most of my years – and my direct boss for many years.

When I started at SAS, they had two 1100/80A systems at maximum config, one running USAS RES, called Resaid at SAS, the other one split in a production partition for Check-In and their old Cargo system, running in 494-emulation mode, plus a test partition. We soon got an 1100/60 system for testing [and some network application] and the full 1100/84B got dedicated to production. At that time SAS was running the systems at absolute maximum, most days with a utilization over 90%. On peak days certain “heavy” functions had to be disabled just to carry the load. Management had a problem accepting that the systems were not big enough and called in numerous outside “independent consultants” to analyze this, wasting lots of good money. Most of them gave up, but at least one team frankly said, that according to their theories the systems should not even be able to do the work they already did [at that time IBM said that systems should never be run at more than 70% utilization].

The new system that could help SAS was the 1100/90, but SAS management was very reluctant to buy one – and had in their minds a strategy to move it all to IBM. In one meeting then head of SAS Data, Curt Ekström, said that buying an 1100/90 would be “over his dead body” – but a few months later they did buy an 1100/93. The hardware installation was a challenge as this was the first 3-cabinet version of this water-cooled system and the water flow was complicated; when they first turned on the cooling water it resulted in water under pressure colliding inside the complex – and cascades of water streaming from the system! This implementation also included moving to a new EXEC and the move to IRV3 [Integrated Recovery version 3, a merge of version 1 (TIP only) and version 2 (UDS and MCB)]. To test this out we had Ted Lambton onsite for a long period and he and I tested IRV3 in shifts, as SAS did not have test capacity to do this during the day. So, the systems were split each night and Ted worked each evening and then I came in at 4 AM each day and continued testing until the load required all parts of the system being used for production – and I did my normal work during the day.



The idea to move to IBM remained high on the SAS Data management wish list and they launched a project called Terese. This moved a long way involving many employees and tons of consultants: but after years of work not really getting anywhere the project was stopped – after an estimated cost of 800M DKK! – a sum quoted in a Danish newspaper article [which probably also included the loss SAS made when dropping out of the Amadeus development cooperation]

My main work at SAS was throughout all the years to maintain their massive amount of EXEC local code and to handle all problems across all OS2200 products. When I started, they had over 10,000 lines of code on EXEC 38R5A, but over the years I removed a lot of this, to be under 3,000 lines.

Some of the code became obsolete [like Message Recovery for COMPOOL and CMS7], whereas other got implemented as standard code or we found a way to use standard tools. Another way to reduce the code was to move EXEC local functionality out of the EXEC, such as replacing local TP key-ins with key-ins controlled by a user program (QIMON). Since the systems were running close to capacity, we also hit a fair number of problems, and my ability to analyze dumps, learned at RECKU, got refined while at SAS. Also due to lots of testing I got a lot of practical experience in operating the systems.

One other development was the story of TTL, HVSTAT and TPM. TTL [transaction logging] was developed between SAS and LH in the early USAS years and then carried to the USAS EXEC group [Dennis Konrad, Roland Angvall, and others] and released to the airlines as part of the ASP (Airline Support Package). It was decided to integrate this code into standard EXEC, in 39R7, and in doing so Roseville changed the control of TTL somewhat and renamed it HVSTAT. Over the years I moved the TTL code at SAS in the direction of HVSTAT, but when reporting some problems with HVSTAT I learned about the new TPM package being developed. This resulted in a long period of me testing TPM [in a wonderful cooperation with people in Roseville] and the migration of all TTL code to TPM, including an interface module to the many TTL analysis routines at SAS. I later had another good project at IBERIA to adapt this interface to their (little different) use of TTL. This also gave me a head start on TPM and deep knowledge into this EXEC feature and made me develop some analysis programs I carry around the world.

Another special task that I have carried around was the development of TRMN8R (the Terminator). No matter how well balanced and well developed a transaction system may be there will be situations where processing stalls for one reason or another. The way to handle this has been to call a system analyst to look at the situation – in many cases resulting in a reboot of the system. Over the years certain procedures for analyzing these situations were developed and the idea arose and was defined by Erik Andreassen at SAS to have a program do what system analysts would typically do. The result became the TRMN8R program, developed to SAS specifications, which would capture system information and through a script language allow an analyst to define the proper actions. This became a great success and avoided many system reboots in the years to come. Since TRMN8R was partly SAS property and required royalties to SAS for others to use it, I have developed other systems to do this type of analysis, such as TXQCK for LMS in Eagan.

A little different project was the USAS SYS upgrade to 11R2 in 1993-1995. SAS was running with the original USAS applications installed in the late 1970s, generally based on level 6R3. With new technologies like Open-OLTP it became desirable to upgrade the core application

USAS SYS [and with-it USAS AEI] to the latest version 11R2. Most of the work was done by USAS people and a few USAS experts came to SAS from Eagan, and other places. But since the initial set up required several system-oriented tasks I got involved and had the pleasure of doing all the initial USAS SYS gens and set up and adjust ALTENP to SAS standards. I also got involved with developing some new USAS functionality, UB bank, deferred MCB messages, so I got quite a bit of USAS knowledge.



USAS SYS 11R2 was installed in AG 2 in parallel with the old running USAS SYS. I was also deeply involved with the actual cutover process – when we finally made it. It was a long and tuff project, amazingly run by Ingemar Jønsson, and it was in danger of being shut down several times along the way. But we made it and Lars Bøggild and the other USAS SYS guys have since made sure to keep their USAS SYS up-to-date – and in many cases way ahead of the standard and contributing significantly to the continued development of USAS SYS, mostly in the area of Extended/Mixed Mode support and with Lars' own debugging tool VDA.

The USAS SYS 11R2 project also gave a few international friends with all the USAS consultants involved. The project started with Bert Barker and Dave Brandsness, but Bert was soon replaced by James Stramel. Later, Jürgen Voight and Gary Sowada became the key members – and they were the ones that carried it through in the end.

A side project to USAS SYS was the challenge to make TTS work in the SAS environment. SAS terminal standards were based on an old Ericsson Alfascop terminal, whose protocol is very much like the Uni-scope – but with some significant changes [e. g. different supervisory sequence]. This required major changes to TTS; so I decided to bring it under COMUS control for most of which I got my old friend Tore Ask onsite.

For many years SAS continued to be “one of the big OS2200 sites”, for many years running at the bleeding edge of Unisys OS2200 technology. The 1100/90 was replaced by the second 2200/600 in the world [and the only one on revenue in 1994 – installed 15th December and online 28<sup>th</sup> December!]. The 2200/600 gave me the opportunity to participate in 2200/600 system testing, running around the clock in Eagan, in teams; but this mighty machine just worked – the central complex at SAS had only one failure in its lifetime, quite an achievement in those days.

The next system should have been the first IX4802-5 in the world. I was deeply involved in the testing of this system (see story under Roseville), but in the end the larger IX4800 systems got delayed, so SAS took temporary delivery of a 2200/900 [water-cooled like the 1100/90 – so fortunately we still had the plumping underneath the false floor]. SAS also installed numerous test systems running on 2200/500s and the production systems grew into IX5800 and IX6800s. Later ClearPath Plus system were installed and when they shut down, they had two Dorado390s with 16 IPs each [installed at two different locations], both split in three partitions, and acting as DR for each other.

In all the years there we performed many software updates and new software implementations. With SAS running at the edge of development I often got involved with the newest software and for years we were constantly running Beta testing of new software releases. One advantage to do this was that SAS at some point defined a need to have a “destructive test system (TEST-X)”, where we could do tests that you would normally not dare do on a system with many users on. We also did special performance projects, e. g. leading to the “TIP Fastpath code”, developed onsite at SAS by Roland Angvall, Pete Day, and Jan Murelius. Another exciting project was the definition of XTC – or Big Mipper in those days (more about that later).

Most of the projects and implementations at SAS were successful. But over the year we also had a few “disasters”, the biggest being:

- SAS260: One dark evening we lost both legs of parts of a database file, as a system support person sitting upstairs was doing an on-the-fly move of some files while an engineer in the machine room asked the operators to take a disk out for maintenance. The reaction was instantaneous aborts, and a partial long recovery was started.



But afterwards application people found out that some data was missing and after some analysis we found that I/O errors to TIP files during commit processing occurring after a program had terminated was not reported anywhere. This led to a long process to close this hole – and cost me three trips to Roseville. SAS260 is the internal UCF number of this incident.

- XPC 11hours: SAS was in the process of moving from the 2200/900 to the IX4800 and during this period we had an XPC (global cache) connected to both systems. We had been testing with this for a long time without any issues. But one night during operator training on the IX4800 a fatal error message was issued on the 2200/900 console. The operators called me – and I immediately told them that they only had one solution: system-wide Long Recovery! Unfortunately, the dumps were 23½ hours old, so this was absolutely a worst-case – and it took 11 hours before the system was up again, just before 7AM next morning. SAS treated this incident as if it were a plane crash and called in a committee. We found that somewhere a well-hidden note said that some microcode update was needed to run the XPC connected to both systems. We never got any serious complaints about this – apart from that I did not have the authority to initiate such a big operation; but they also concluded that it was good that I did it as that brought the system up in time for next day - and it was the right decision.

Running at the edge of capacity and, at some point, one of the highest transaction rates we had quite several problems, in many areas at the OS2200 product suites. Being the one (or only few) Unisys person on-site meant that you got involved with all of these and learned that there is only one thing to do: dig into it! You cannot just sit back and say, “I have not been trained in that”. My primary area was [and is] the EXEC and the dump analysis skills I learned at RECKU really got refined at SAS. I had a procedure for the dumps, which meant I had an overview analysis of an EXEC dump available within 10-15 minutes after the stop. We often developed a fix for the problem before it was submitted as a problem. Or the few times, where we used “dirty tricks” to overcome a problem. One of these clearly impressed people in Roseville – enough to quote in a PLE that they could call me [Archived PLE 17023993].

Soon after I started at SAS the Unisys software support staff was greatly reduced as SAS did not continue with the Unisys-provided USAS developers. For many years I continued to have one more person working with me on OS2200 support, sitting with me onsite at SAS; I therefore had the privilege to work with great people like Anders Öhman, Mikael Rattenborg, Mark Baltzersen, Jan Lauridsen, and Niels Holm Sørensen.

I was assigned to SAS for the most of 18 years – but there were some breaks, where SAS Data management [or whoever] decided there was no need to have an [expensive] full-time consultant from Unisys on-site. The first break was in 1995 where I was assigned to Telia instead. While I was away at Telia, SAS decided to try on their own to upgrade their software, a requirement to handle year 1996 in OS2200. For that they hired Steen Jarbøel, ex-RECKU and ex-Unisys, as a cheaper consultant; it did not work and so they called in Unisys and Tore Ask and Palle Hansen came onsite; it still did not work. So, in October SAS gave up and signed another contract for me; I then had only two months to make the system upgrade, a project that normally took 6-12 months. I had some very, very long hours, but we made it.

My long assignment ended in 2001 – very abruptly. As usual the contract had expired while they were still negotiating, but one day Per Voss called and said it was not moving and asked me to leave SAS immediately. I very quickly told the SAS staff what was going on and walked out the door.

I came back to SAS a few times later for month-long assignments. We did a couple of software upgrades and brought in their last OS2200 systems, two 16x Dorado390s. It was great to come back, and it quickly felt like old days; those were wonderful visits. I also felt that the staff at SAS enjoyed these periods of hectic activity during upgrades.

I also must mention and extend a great thank to the staff at SAS Data – my colleagues and friends for many years, more than most people at Unisys. The core people were the XK-P/VU-P system support group: Erik Andreassen, Palle Emil Larsen, Dan Sjøland, Jeanne Eismark, Kurt Jensen, Niels Stoffregen, Holger Sønnerup, Robert Jensen, Bo Simonsen, Inge Westergaard and Anne Hertz. There were also the USAS SYS people: Ingemar Jønsson, Lars Bøggild, Allan Thompson, Peter Sørensen, Finn Rasmussen, Ole Nielsen – and then all the USAS application folks, as well as other system support people: Egan Larsen, Sune Bjerregaard, Birger Brandt. Finally, I must mention all the operators with whom I have spent so many hours, mostly testing, but also bringing a system back up again: Birthe Andersen, Lasse Meinertsen, Preben Nyeland, Peter Glynn, Torben Rolsted, Alexander Mauritzen and the “young team” of Maybritt Therkelsen, Randy, Birthe, Bjarne Frausing.

I had a great time at SAS, and I am pleased and proud of what I did there. SAS is deep in my heart, and I will always feel an attachment to the airline. I will probably always try to book SAS when I fly – despite being very unsatisfied with the lack of in-flight service they provide today. Furthermore, the friendship with the people there never stopped –and I hope it never will.

On midnight between 31<sup>st</sup> May and 1st June 2013 the OS2200 systems at SAS were finally powered down. I could not be there that night [I really would if I could], but I attended a lunch on the day that the last of the OS2200 SAS people had their last day at work, on 31<sup>st</sup> July 2013. There was a strange and sad feeling in the air at that lunch – the end of a fantastic and outstanding era.

As noted, then during the many years onsite at SAS I had a few breaks with shorter term assignments at other large sites in Europe, so the story continues with these sites.

## LUFTHANSA

The German airline Lufthansa was probably the biggest OS2200 site in Europe at the end of the 1980s. Consequently, they were the first site in need of implementing the Big Mipper feature (or XTC as it was called later), where multiple OS2200 systems could run in parallel with access to a shared database. The project had been underway for several years, but in 1991, it just had to be completed. At that point, the OS2200 software support group at Lufthansa was run by my ex-SAS colleague Mikael Rattenborg, and he asked for me to come down to Frankfurt to help with the final testing and verification before going online.

At first SAS would not let me go, but when we arranged for one special terminal in a SAS back office in Frankfurt airport to have full demand access to all SAS systems, then SAS accepted that I could go – and in emergency situations I could help debugging problems at SAS from this special terminal. This was in the days before the Internet and terminal connections were dedicated and directly [possibly via multiplexed lines] connected to the systems.

I moved to Lufthansa in May 1991 and started working on the final testing as well as migration procedures, together with the Unisys staff at Lufthansa [Mikael, Jürgen Brors and Torsten Klebe on OS2200, Ulrich Mueh and others on TELCON and a great CE staff – as well as Lufthansa support people, like Dieter Lothar].

The work was very interesting – and challenging. I had been involved in the XTC development [or Big Mipper as it was first called] right from the beginning - described under Roseville later.



After years of testing and evaluation it was time to bring XTC into production and I was asked to help with the final verification and check out procedures. This called for numerous test sessions and for considerable writing of documents – fortunately in English. We also had many meetings, and it was good for me to brush up my German by accepting that meetings were held in German [but I still spoke English if I had something to say].

The implementation had multiple phases – and we got off with a very bad start! First step was to introduce an EXEC with the XTC features integrated, but without the hardware locking module, the RLP, (**R**ecord **L**ock **P**rocessor) or any shared disk. The EXEC was booted on the production system one early morning, but it soon became clear, that something was very wrong – and it was realized that the database was randomly corrupted. Now Lufthansa showed their German efficiency and immediately ordered a full Long Recovery of the database up to the point where the new EXEC was booted; they started multiple runs in parallel and it just worked! We later found out that the problem was caused by all deferred updates from the Rollback Page file were written over four records instead of one [or every fourth record, or something like that] – because XTC allowed up to four parallel systems [hosts].

Towards the end and for the cutover we also had John Carter from Roseville onsite. John is a fantastic guy, one of the best I ever met; John was also a big guy. This gave him a challenge with the access ports to the computer room. They had two doors and a small room in between – and with a scale in the floor.

You entered your access card and your code then the first door would open, you stepped in, and then the door closed behind you; then hopefully the second door would open, and you could walk through. The idea of the scale was probably to detect if you tried to pass two people through at the same time – and it must have been set at 150kg (John's weight). Because if John was carrying a listing, manual or whatever the second door would not open! So, when we went to the machine room, we always had to carry things for John. They had the same doors for entering the building, but they were clearly not so sensitive, as he could go through those with his briefcase.

The actual cutover on 22<sup>nd</sup> January 1992 was a great success. All other steps in the migration had worked fine and we now had a 2200/644 running with both shared and local disks and locking using the RLP. For the cutover we had to add a 2200/622 to make it a two-host XTC system. It just worked! We had a safe margin of planned outage time – so when everything was running and fine and we still had scheduled time then Dietmar from Lufthansa got so excited that he decided to try to reboot the original 2200/644 – and it also worked perfectly!

I have been back to Lufthansa a few times since then. One trip was a Check-In system outage where I was called up in the Unisys office in Denmark and asked to get on a plane to Frankfurt ASAP. This is the only time I have booked a plane ticket from the taxi taking me from the office and home to pack the minimal – and directly on to the airport. When approaching the airport, the captain announced a delay due to severe problems in Frankfurt airport – I knew why! When I finally came to the office in Kelsterbach the system was running, but there was a major problem analysis task. I continued to maintain contact to the Unisys Lufthansa team until the end.

And then in October 2018 I received a mail saying that the last OS2200 system at Lufthansa had been powered down, again an end to a fantastic era.

## TELIA

In 1995 SAS for the first time decided that it was an unnecessary expense to have a full-time Unisys consultant on-site. So instead, I was assigned to the Telia account in Sweden, the largest OS2200 installation in Scandinavia.

I cannot remember why they needed an additional OS2200 resource, but for four months I flew to Kalmar [often via Växjö] each week. I shared an office with Leeni Heinonen, but most of my work there was done together with Ingvar Nilsson, their EXEC and basic OS2200 software guru. I got involved in their daily operations and system upgrade planning, but with talented people around I did not have a lot to contribute. But then one night, where I happened to be onsite, their test system TTST suddenly ended up in a critical situation that required a JK13 boot. It did not work – and I got deeply involved and managed to get the system up overnight; I believe I earned some credit that night – and initiated some changes in their procedures.

This support assignment abruptly ended in October 1995 when I first had to go to Beijing and CAAC [for the first time, see below under TravelSky] and later I was rushed back to SAS to handle their critical year-1996 challenge [see under SAS].

Sometime in early 2000s Telia decided to outsource their IT to Tieto-Enator; but as they had no OS2200 experience they selected to sub-outsource the OS2200 operation to Unisys. The OS2200 systems in Kalmar were moved to the Unisys Outsourcing Centre in Alingsås [a great effort, masterminded by Jan Murelius]. 17 key Unisys-oriented Telia people were employed by Unisys and many Telia OS200 developers employed by Tieto-Enator. In 2004 one key person, John Burgess, responsible for the critical system middleware PRODLIB, was scheduled for retirement and the Unisys Telia group was looking for someone who could take over the responsibility for this critical program environment, of which a lot was written in MASM (the OS2200 Assembler); and so, I found myself assigned to the Telia project again.

For the next year I travelled, both to Kalmar [where most of the support people are] and to Stockholm [where John Burgess and Per-Arne Nyström were] first in Haninge [or Handen, as it was renamed] and then to Älvsjö [the Tieto-Enator office]. This gave me some nice revisits to Kalmar and good visits to Stockholm, especially Södermalm.

The PRODLIB software is impressive. It was developed in the early 1970s by some talented people and contains features that were to become available from Unisys many years later [and some not yet]. One such is the screen- and program-handling BGEN, packed into one component; this allows you to define a transaction input screen and process it without a single line of coding!

I probably reached 50% of fully understanding of PRODLIB by the time that John was scheduled to retire, in December 2005. But John never stopped as he was employed by Unisys as a part-time consultant; this was great for Unisys and for Telia; but it meant that I never got to be fully confident in handling PRODLIB on my own.

A few years later I was contacted by Per Anders Wicklund at SAS. Per Anders was another guru at Telia in 1995 and single-handedly ran a small Telia Centre in Malmö. He later moved to SAS, where he soon became a USAS RES guru, taking over from Leif Norst. When SAS started to go away from OS2200 and USAS he wanted something else, and in the end, we got Per Anders employed at Unisys and assigned to the Telia project, where he knew the systems in and out. Per Anders got assigned to PRODLIB and with his talents he soon took it over completely, so I had even less to do on PRODLIB. In July 2013 John finally retired completely and PRODLIB was now in the hands of Per Anders, with me in the background as a backup.

I have continued this assignment until mid-2018, where some network changes at Telia meant that I could not access Telia at the same time as accessing Unisys or other systems. To do that you needed multiple instances on the PC, and I did not want to add this complication, so we decided to stop my support of Telia, which had already been limited over the last years. The Unisys systems at Telia was finally shut down in June 2019.





## AMADEUS

In 1999 I had another break in my assignment at SAS and so I ended up in the support group at Amadeus in Erding, outside Munich in Germany. Amadeus is the giant airline booking system, originally started by Lufthansa, Air France, and Iberia with the purpose of having one common reservation platform for all travel services, a front-end for all the airline back-end inventory system [controlling the individual airlines' own fleet of airplanes]. They needed a computer system, and as all parties were Unisys-airlines, Unisys [as far as I have been told] somehow made SAS join, assuming this would improve the chance of Amadeus choosing a Unisys system. How wrong! Amadeus selected an IBM TPF system [probably because Unisys did not have a system big enough – only in design].

But Unisys got a piece of the cake as vendor for Amadeus' Fare Quote system. Fare quoting used to be a simple application updated from a central source once a month, but in the highly de-regulated market of the 1990s Fare Quote became more and more a competitive tool and prices became a jungle, nearly updateable in real-time [and unfortunately it still is!] So, the small bite of Fare Quote became one of the biggest Unisys OS2200 systems ever. It had a central core of a 4-host XTC system, but they were surrounded by numerous satellite systems, some OS2200, but also many [and growing number of] IBM RS/6000 RISC systems. The satellites interfaced with the central system through a giant memory-to-memory caching interface [designed by Paul Cameron].

Some key people had left Amadeus, so they needed a hand. I was only there for three months at the end of 1999, but it was another encounter with a great group of people: Steve Chisholm, Eric Staubinger, Charles Kendall [he ran mountain-marathons – how can you?] and Danny Doyle. My major task there was to get TPM working in their environment.

## HIS/LMS/DIGISTICS

When my assignment at SAS was over, I needed something else to do. At that time Unisys was trying to sell our LMS Cargo outsourcing service to SAS Cargo and that way I learned that the HIS (Hosting and Integration Service, running the LMS system) could use an additional operation support person. So, in 2002 I joined the HIS/LMS support group as a remote support and I could then come to Eagan 2-3 times a year, as Gary Moerke said – never happened! We managed to sell the solution to SAS Cargo, and I also got involved in the migration for SAS - replacing their old 494-emulated application.

The SAS Cargo migration project brought me to Eagan for three weeks, which was very good for me to get to know my new colleagues, at that time primarily Kirk Nosbusch and Mike Stella. They have both now retired, so I rely a lot on Steve Grotbo and Iqbal Bedi - and previously Gary Sowada, all of whom I fortunately have met in person – it really helps to “know” your colleagues. I still have this assignment and act remotely in a group of six to eight people providing first line operational support for this important and growing application. We make great use of the time differences, so for many years I carried the HIS Operations pager, via my cell phone, whenever I was at work; this meant that I took care of the pager while the US-based people (rest of the group then) could sleep without being disturbed (1-9 AM). Lately people in India (first Kiran and Aditi, then Neethu and now Vasudha) were added to the support group, so now they mostly take care of the US night hours. Having been part of this for 19 years I believe I can handle a lot of tasks, even though it would be nice to go to Eagan from time to time to be closer to “my work” and the people I work with - albeit more and more of them now work remotely too.

The SAS Cargo migration project also gave me a brief insight into the fascinating world of Cargo. I attended a lot of project meetings at SAS Cargo in Copenhagen, and it was amazing to see how our people, mainly John Huber, Tom Zurick, and Tony Robertson who knew the Cargo world.

They and were able to explain everything that SAS Cargo brought up as questions and concerns. I also got some exciting tours of the SAS Cargo facilities: a parcel requires a lot more logic and handling than passengers – they cannot just go to a gate when asked!

I believe I have made two major contributions to the LMS system: OS2200-side of EABS system and TXQCK transaction monitor. The EABS system controls the release of all code changes to the systems, all performed on-the-fly. There was an old system, mostly written by Elmo Spielman, but it was quite cumbersome and required several manual steps – and it did not have a rollback facility. So, I rewrote the OS2200-part and tried to satisfy all the requirements. It has a server part, written by Mike Stella, where all changes are registered, documented, and all levels of approval are kept; once a change must be implemented the server initiates a process on the 2200 systems – and that is where my [rather large SSG] routines take over. The OS2200-part of EABS assures that all programs are updated across the nine systems used and that status is returned to the EABS server and required emails are issued. I think this software control and release mechanism is superior to anything else I have seen.

The LMS system has over the years suffered from numerous situations of transaction queuing and that was the major source of paging of HIS Operational staff [often many times per shift]. Since I could not use TRMN8R, developed for SAS and requiring royalties to SAS, a new, simpler transaction monitor system was developed [using my QIMON program and SSG skeletons, i.e. very flexible]. It was named TXQCK and it has saved many people from being woken up in the middle of the night! It also reacts a lot faster than paging and user analysis – and does it in a consistent way. TXQCK could probably be very useful in most transaction-oriented sites.

When the LMS solution was sold to TravelSky, both in 2008 and 2013. I was also involved in the set up in Beijing – since I happen to have a leg in both camps. Since only a few HIS people has demand access to the TravelSky systems I also enhanced the EABS system to a point where all processes could be initiated by a simple mail – but it was never really used.

My work for HIS/LMS is a little different from what I have been doing in other assignments. I have mostly worked at the deepest level of our software [EXEC, basic software] and generally been able to even boot the systems. Here this is handled, very well, by the Eagan support group: Jay Stevens, Darrell Anderson, Caryn Dunkel and before Sharon Loftsgaarden, Dennis Bowes, and Dick Arten. At very tricky situation [like a disk issue not long ago] I really miss opening the console to see what is going on. My tasks are at a higher level, but here I take advantage of my limited USAS knowledge – and it seems to be good enough. But I miss a lot of knowledge into newer technologies like MQ and everything in the surrounding Windows [and Linux] systems.

Now at the end of my Unisys years, this is still my main “daily” task, although the hours used here are also being reduced.

## IBERIA

IBERIA has been a long-time user of Unisys systems, initially running 494s like SAS. In the 1970s they moved to the USAS based airline system alongside SAS and Lufthansa and consequently the IBERIA system looks very much like SAS. On top of that then the system software at IBERIA has since 1990 been handled by Anders Öhman, until he retired in 2014, who “grew up” with USAS systems together with me at SAS. Since Anders and I are also good personal friends we have for years been doing things in parallel and interchanged knowledge and experiences.

Over the years I have been to IBERIA numerous times for different meetings and consultancy as their long-term key manager Christer Henmark often wanted “another opinion or view”.



So, whenever something major happened and some deep analysis was required, I was asked to come down to Madrid for a few days and work with Anders and other people to “sort things out”. But I have also had other projects and educational tasks at IBERIA.

My first longer-term project at IBERIA was to re-do what I did at SAS: convert the IBERIA TTL routines [rather different from SAS versions] to work with TPM. I could use the same basic concept as at SAS, but the volume at IBERIA was much higher and was a challenge. I had three good weeks at IBERIA – also meant that I became a little closer to the IBERIA staff. In connection with this task, I also acted as vacation stand-in for Anders and hopefully built some trust that someone else than Anders can help them in critical situations. I later also updated their TPM interface routine to handle the DSTA-version of TPM.

Following a serious fire at the IBERIA computer Centre [took out their XPC, fantastically handled and recovered by Anders] they became quite serious about Data Recovery capabilities. Consequently, IBERIA has a complete duplicate of their main system at a remote location, a cold standby. As this is a bit of waste and rather costly IBERIA would like to use this spare capacity – and then the idea of using XTC came to mind. I did a major analysis in 2008 and produced a report for an XTC implementation; but so far it has not happened, primarily because IBERIA wants to split the two XPC-L engines between their two sites, and they are too far apart. The XTC idea has popped up at few times since then, lately when they wanted to split off Availability Manager logic in a separate host [would make the communication side of XTC much simpler].

In 2012 Anders decided to retire and a desperate need for a replacement came up. Fernando Nicas stepped into the role, but he and IBERIA felt he missed some experience to handle certain critical situations. The initial solution was that Anders would still be available 25% and that Jan Murelius and I should come down for onsite support approximately one week per month.

So, in October 2012 I started my “commuting” to Madrid to support IBERIA and train Fernando. But one year later, Jan Murelius also retired, and Fernando had proven that he could handle nearly everything, so the support has been reduced – and Anders was still available when needed.

In 2014 Unisys signed a major support agreement with IBERIA and the onsite Unisys staff grew considerably, gradually taking over tasks from the retiring IBERIA staff. Unisys also took over responsibility for operating our systems – in addition to an already quite considerable USAS staff. But as Unisys Spain had to lay off many people the first ones were the semiretired, so that meant that Anders was laid off and no longer available. So now IBERIA again needed external support for critical situations and Jan E. Jensen was added to the support staff, providing onsite support, along with me, when needed. The Unisys staff at IBERIA was expanded with several non-OS2200 people, so in 2015-2016 Jan and I carried out many training classes to teach these new people OS2200 – with a growing number of people attending remotely via Lync while recording the sessions.

## TravelSky

I first went to TravelSky in October 1995 (then CAAC – **C**ivil **A**viation **A**uthority of **C**hina, a semi-governmental organization). I was at that point working at Telia and could not release my passport, so it was decided that I should go via Hong Kong to get a visa. That turned out to be difficult and I was unable to get a direct flight, SAS had a direct flight to Hong Kong, or even a two-leg flight, despite having an unrestricted C-class [business class] ticket [I think the last one I ever had! So, I flew SAS to Paris and then Thai via Bangkok to Hong Kong. This was my first Asian trip – and then with Thai business class; but I did not find it so outstanding.

My return trip was fortunately with SAS non-stop – and their business class was wonderful, my first meeting with Remy Martin X.O. It took one business day to get the visa, so I had one day in Hong Kong. I had a meeting with Ed Tuomy, then remotely managing CAAC, but then I had the rest of the day off for sightseeing in this fascinating city.

The CAAC office was at that time on Chang An Avenue, the huge main street, running through Beijing from East to West. CAAC had just bought a 2200/700 [a 2200/500 with a “built-in” XPC, the system-global data cache]. The XPC was new to them, and they needed someone to tell them about it and test it.

My primary task was to update the operational procedures for handling the new XPC-logic. This affected all part of file handling and back up, procedures for I/O-related operation, planned stops and all type of boots. CAAC also uses the RCVRY package and some of the programs needed upgrading. The Chinese are very, very careful with their testing, so I participated in numerous test sessions. On top of this I also spent some time with their primary OS support person Zhu Hiaoxing reviewing and updating some of their local code for the EXEC and other products.

For many years the main external support person for TravelSky was Garth McClellan. But in 2004 Garth suddenly decided to retire, so they missed a person to maintain their EXEC local code, provide operational support and act as third-level support. Garth recommended me for the job and since they had seen me before (I must have done OK in 1995) I was asked to take over this responsibility. This is a wonderful assignment, and from time to time rather challenging.

My first trip to Beijing after taking over the support responsibility came in February 2005 – a very long trip, kicking off the XTC project. From the moment I arrived this trip was so much different from the trip in 1995 – it was amazing how Beijing [and China] had developed in 10 years.

On the first trip I worked with Ray Salisbury, Zafer Erbas, and the Project Manager Debu Mahapatra together with a lot of TravelSky people to define the requirements in bringing XTC into the TravelSky environment. Many other people were involved over the next three years, but with Debu, S.K. Jain, Unisys onsite USAS support, and Rob Jamieson (TCIS APAC) as the key resources. Just around the Olympics Darius Engineer took over as Project Manager – and he has been there ever since.

The XTC project resulted in numerous trips to Beijing over the next years. To make a long story short: we made it, and well in time of the 2008 Beijing Olympics the TravelSky systems were running as clustered XTC systems with sufficient capacity to handle even extreme loads. There were numerous major cutovers, where one of them was such a big event that “everybody” had to attend. There was a rope across the operations room and only the critical staff was allowed inside, but many people stood behind the line. There was also a big press presence, and many pictures were taken. When it was well completed the CEO of TravelSky came down to thank us – and it was quite a surprise to find that it was Zhu Hiaoxing, the system analyst I had worked with in 1995 [we both immediately recognized each other]; it is not often in the Western world that a System Analyst becomes CEO of a large corporation.

I was also requested and paid for by TravelSky to be onsite, both for the Olympics and the Paralympics. So, I spent six weeks of the fall of 2008 in Beijing.

The Unisys Olympics support team was installed at Crowne Plaza at Wangfujing [at approx. 5000RMB per night, approx. seven times our normal price]. On opening night, the hotel threw a big party for all the residents, so we watched the Opening Ceremony on a large-screen TV in a bar area – with the fireworks happening outside our windows at the same time.



After the Olympics there was a quiet period at TravelSky, but then the natural growth of business in China overtook the extra resources for the Olympics and it was time to expand the systems again. All systems were upgraded to four-host systems and newer hardware was added. We also went through a series of software upgrades and generally TravelSky has requested me to be onsite for the first production system software upgrade [often the CRS system] and for the upgrade of DCS, the most critical system. I like the work at TravelSky and believe I have given them a level of confidence in the support Unisys can provide.

Working at TravelSky is not the easiest. First, they of course speak Chinese – and I do not! Most of the people in the OS group I worked with spoke English well enough for technical discussions, but they are clearly uncomfortable with it and speak Chinese between themselves when possible.

This means that I often sat rather alone in the office without any idea what is going on around me; I have had situations where there has been an issue on a system and yet they work it themselves without asking or involving me. Talking about sitting, that is also a different experience. First, they do not have special guest places, so they just place you where someone is out on that day – and next day when you walk in “your desk” is taken. The other specialty is that you often sit quite low; some chairs do not work too well and when you put a big foreigner in the seat the pumps cannot hold, and you sink to the lowest position; then you must lift your arms to the level of your breast to reach your keyboard.

The TravelSky staff are not very social. There have been quite a few times where they have taken frustrating long time for even rather simple OS2200 tasks – at least in my view. And then suddenly, without warning and information, they get things done. Contrary to what some people expect then they have quite normal work hours and are away on vacations – and some of their national holidays mean that the office is empty for days – but there is always a member of the OS group at work during the day.

TravelSky is probably one of the largest OS2200 systems in the world – and any critical migration there calls for high level attention within Unisys. Roseville is always ready on the sideline when we are doing things at TravelSky. They are also often at the bleeding edge of our technology and quite a few features have been developed to satisfy special requirements at TravelSky. Over the years the ICS and CRS systems grew into six-hosts systems, and in 2016 the main site, Dong Si, has one Dorado390, four Dorado790s and one Dorado890s [all maximum 32x]. The DR/test site, San Li Tun has two Dorado280s, one Dorado380, one Dorado880 and one Dorado6380 – plus XPC-Ls to act as DR and make some test systems true XTC. The production systems run at peak an accumulated load of up to 9,000 transactions per second! In 2018 a project was started to move the OS2200 systems to the new TravelSky facilities in Hou Sha Yu near the airport, introducing new intel-based 8490-systems. Hou Sha Yu is a massive place, nearly a city on its own with four nine-story office blocks, a separate data Centre with 32 machine rooms, a cafeteria building, and three tall apartment blocks for employees on shorter or longer visits. The latest move of the DCS system in May 2019 was also a giant event, with hundreds [including highest level management] of people present, all run from their impressive main operational Centre (ESS). The DR site is now at Jia Xing near Shanghai.

My key roles at TravelSky were two: to maintain their EXEC local code, including recommendation of PCRs for each new EXEC level and general help with EXECs – and then as onsite “expert” during critical system migrations. The last role means that they call me onsite whenever they need to move critical systems, must switch XPC-L [which means a total application outage] or implement new software levels. There is nearly always some small [or sometimes larger and more critical] detail that I need to handle and thereby justifies the expense of getting me onsite. So, I have spent quite a few hours in the TravelSky machine room, where access is very restricted.



It is amazing how often I have been put to work during TravelSky upgrades or system moves. The OS2200 software stack is generally very, very reliable, but it looks like the TravelSky environment is so special [large or in other ways different] that we hit problems not seen anywhere else. We always have had Roseville on stand-by for major system changes and many times we have ended up with me sitting in the machine room analyzing dumps, passing details to Roseville people, long conference calls – and heated discussions.

But so far, I have been able to carry it through, which some people, TravelSky being one part, consider to be very valuable and a reason for constantly calling me back for these occasions. It has been challenging, it has been very tiring – but also very rewarding to pull it through again and again.

The Unisys teams coming to Beijing were a very social group and we do a lot together. In the beginning the key person was Debu Mahapatra, but after the Olympics Darius Engineer has taken over the role as the key project manager. There are numerous other project managers for special projects – and for my OS2200 support work the key person is Rob (Robert) Jamieson, responsible for ClearPath services in Asia. It has not always been clear what grants you access to the lounge in our hotel [apart from paying a premium rate], but if Darius or Rob are there being their guests work.

Talking about taxis I must mention one other thing. The first years I always had a hotel car or a Unisys-provided chauffeur to pick me up in the airport. Nowadays I feel confident enough to take normal taxis [approximately half price]. When I used to stay at Prime hotel, I managed to say the name of the hotel in Chinese and then I was of course taken the shortest and direct route to the hotel. I have not yet learned to pronounce Crowne Plaza in Chinese, and there are many of these in Beijing, so I rely on the hotel business card to get me there.

## SHORT-TERM CUSTOMER SITES

I have been involved in numerous other smaller projects over the years, but they all somehow are part of my Unisys life. Some of these go back to my early years in Unisys, whereas others are more recent.

### Fjerndata

Fjerndata was an IT service bureau in Norway and one of the biggest Nordic OS1100 installations in the late 1970s. They needed extra hands, and I was assigned to provide onsite support, together with Ian Hepburn and probably others.

Like many of the early OS1100 sites they had some brilliant people, and at Fjerndata the star was Øistein Nilsson [I still carry his file of great utilities, all written in MASM]. Unfortunately, I do not know what happened to the site and the people, but all is gone years ago. During one of these support trips, I had to support the site all on my own.

### Norsk Hydro

The Norsk Hydro site was in Porsgrunn, Norway, some 150km from Oslo. You could drive to there, but in many cases I took a small domestic plane from Oslo to Skien, quite a flight over the beautiful Norwegian mountains. Norsk Hydro had their own version of a Uni-scope handler, developed by Andy Øverby. My task was to help them revert to the standard Unisys handler [then in the EXEC] and then move them on to CMS. This was not an easy task as Andy's handler had some fancy features that neither EXEC nor CMS could provide.



### University of Tampere

In the late 1970s Unisys Finland signed a deal with all the Finnish universities to have each their 2200 systems. Not all universities were pleased with this arrangement and at the University of Tampere they did their utmost to show that the 2200 was not good. They found all kind of ways to make the system crash [not too difficult in those days] and came up with rather special demands. One of these was that they one week asked to have TIP and CMS installed and have it ready by the following Monday morning, a near impossible task. I was asked to help and rushed to Helsinki on the Friday afternoon.

We worked the whole weekend and had numerous challenges, but by Sunday evening late we had TIP and CMS running. On Monday morning the University staff then decided they did not want it anyway and turned it all off again! This “battle” went on for years and eventually ended up in court; Unisys won, but we lost the site anyway.

### Leitz

While being part of the support team in London I only had a few support trips. The first one was in fact three trips to sort out a very tricky problem in CTS Procedures at Leitz in Germany – the company that makes the Leica camera and other high-quality optics. It was a fascinating factory to visit with many old-fashioned facilities yet producing some of the finest optics in the world; the factory was also one of the few places with a Paternoster elevator installed, a constantly running elevator, where you just jump in and out when the elevator went by [they also have one in the Danish parliament].

During my third trip there we also called in two people from CTS development in Roseville to help debug. We were all picked up in Frankfurt airport and taken to Leitz by the German support person – in his very fast Audi 100. It was quite a sight to see the faces of the U.S. people when he drove along at 200+ km/hr., a remarkable speed for persons accustomed to max. 55 mph. We eventually found the problem, but only after I developed a few DAP dump-analysis routines, something CTS never had before.

### Natal Building Society

In London I spent a lot of time testing out Integrated Recovery products and procedures, and that was the reason I got a very special trip, to support one of the first sites to use the full suite of Integrated Recovery: **Natal Building Society (NBC)** in Durban in South Africa. Numerous people were onsite, but things were a mess; thus, one day at lunchtime I was asked to go onsite ASAP. To make it I rushed to the South African Embassy on Trafalgar Square, where NBC had telexed them that I needed a visa immediately. They kept the visa department open just for me, and after a run to Piccadilly Circus and back to get some photos, I got my visa. Then I rushed to the computer Centre to get my tape with all required programs and on to Heathrow, where I met my wife with my packed suitcase – the only time that happened. Maybe this was the reason I ended up two weeks in South Africa without a camera.

We got the mess sorted out and made all the IRV2 components work. I believe one of my major contributions was very non-technical: in an un-used room in this very new computer facility. I used each tile of the elevated floor as a storage facility for each problem and then created a matching problem-matrix map; that way we got control over the massive number of outstanding problems.

This was in 1983 during the period of apartheid and I saw how that horrible system worked at first hand. We had Indian operators – and they were clearly the best that have been able to lift themselves to that level – and they were some of the brightest operators I have ever met.



Established in 1980

Because of apartheid the security at the site was also very high, and the computer room was four levels underground – with two levels of water above to stop of possible bomb!

### Air Canada

Air Canada was for many years a flagship site for Unisys and where TIP was developed, with the systems based in Toronto. But somehow it was decided that the systems had to be moved to Winnipeg and the software be more standardized. This gave a challenge as some interface package was needed to map the old Air Canada application – and support-wise as no one wanted to move from Toronto to Winnipeg. So, Unisys played a major role in this move.

Garth McClellan and the Read-twins, John and David, designed and installed the interface – but there were some performance issues, and I got assigned to look at it. After some initial work with people in Roseville I flew to Winnipeg – another very different place.

### JAT

The Yugoslav airline of JAT was one of the smaller USAS sites, but like the larger sites it required a large staff to tailor the software for JAT's use. Being a small site, it also ran into some problems not seen at the larger sites – and because of the unreliable infrastructure it had to work in (as I later learned). At one time they ran into a problem of being unable to process their Type-B messages and they were piling up. Knut Paul thought I could help, so I was rushed to Belgrade. It truly was a different experience. All the facilities we were accustomed to take for given either did not exist or were very limited – such as proper desks and chairs! Being an outside “expert”, they made very sure I had proper working conditions. The site was in the airport area, so you had to pass numerous well-armed security guards - maybe soldiers - to get to the site, and walking to the canteen you nearly walked alongside the runway. My main contribution was to extract information from a system dump and pass that information on to USAS people – and that way we rather quickly got the problem solved.

We therefore got the processing of the Type-B messages going - many of them totally useless information like the closure of a certain class on a domestic U.S carrier some time way out in the future – but JAT subscribed to this and got it all. We monitored the processing speed and compared to the queue – and then came the remark I will never forget: the JAT staff looked out the window and said: “We think it will stay dry over the weekend, so hopefully the queue will be cleared Monday morning”. Their communication lines were so weak that they were weather-sensitive!

But, for the queue to clear I had to stay over the weekend. It worked – and Monday morning the queue was cleared, so I could go back. This was a very short trip, but so different that I remember lots of it very clearly.

### Lloyds TSB

Trustee Savings Bank has been a Unisys flagship customer for many years and when TSB was taken over by Lloyds; the usual fight between IBM and Unisys systems came up and for years the choice was shifting back and forth. At LBG (Lloyds Banking Group) the Unisys systems kept a major role and thereby grew into one of the largest OS2200 systems. Several bank mergers followed , Halifax, Royal Bank of Scotland, and for the merging of these systems numerous special developments were required, one being DDR, a facility to merge/copy DMS databases on-the-fly. I was not directly involved, however had one visit to their Centre in Wythenshawe outside Manchester in connection with some Audit analysis.



Like most other large systems, a requirement came up to provide built-in system resiliency [at LBG called Active-Active, in OS2200 solved by XTC]. A seminar was arranged in June 2012 to analyze how XTC could be phased in at LBG and I was asked to give an introductory course. Nothing happened for a while, but at the start of 2014 I was back at LBG, first for a Unisys-internal workshop on XTC and its pre-requirement MHFS (**M**ulti **H**ost **F**ile **S**haring)) – a great workshop by the fact that nearly all UK-based OS2200 people were present, later I went there for two courses for LBG staff. Time will show if the system will be implemented – and if I will get involved with this very exciting, challenging site.

## ROSEVILLE AND EAGAN VISITS

I have had the privilege to visit the Unisys facilities in the Twin Cities of Minneapolis and St. Paul many times – and I love to be there, both at Unisys and in the cities.

I will start with the work side. For many years, while onsite at SAS, my manager Bent Strøh arranged that I went to Roseville approximately once a year to keep in touch with the latest developments. Some of those visits were on my own, while others included SAS personnel and special presentations. But here is a list of some of the more special visits:

- First visit! I really had no work to do, but in 1979 during one of the benchmarks in Eagan, Tore Ask took me up to Roseville. He had worked there before and knew many people – and probably had a business reason to go. I will never forget the feeling of respect, when I first stepped into the “holy halls of Roseville”! This was the birthplace of everything I worked with, the place the systems were built, where all designs were made, where the OS2200 developers were and where all problems were solved!
- Big Mipper (XTC) design review. In 1980s some sites, amongst them SAS, were running out of power on the OS2200 systems. Marketing defined this need very simply, with one sentence: “If Northwest is taking over Republic an IBM will have to replace the Unisys systems”. To overcome this Roseville started development of “Big Mipper” which became XTC; where multiple OS2200 systems could be clustered together to give more capacity than one system could provide. When the design was nearly complete Roseville called in the software specialists from the most likely candidates to use this facility, at that point nearly all airlines, for a week-long review meeting. This was a great meeting, getting together with some of the key people from large OS2200 sites like Christer Henmark, Ted Lambton, and some US and NUL people.

We met with all the Big Mipper developers: Phil Hartley, Gary Schwartz, Dennis Konrad, Monica Mauer, now Langsford, John Sjolander, Mike Hill, Albert de Nigris – and more. I dare say that I came best prepared of all; we were given the design documents in advance, and I had spent a long time with Erik Andreassen of SAS going through the documents. So, a few days into the meeting it was: Michael, any comments on this page? Special memorabilia of the meeting were a small rubber Pluto that I had; it was named ARF [a small feature of XTC – and the sound of a dog barking]. For each section I used ARF to signal my feeling about that section’s design: if it was good, I put Pluto’s tail up, if it was bad I put Pluto’s tail between its legs. The XTC team kept ARF/Pluto as a mascot for many years. In the years to follow whenever I came to Roseville, I always spent some hours with the XTC-team to review progress and run some of my tests. Later I had the privilege to install XTC at customer sites, first at Lufthansa and then at TravelSky.

- 2200/600 System Test. This test took place in Eagan but was run by Roseville. Since SAS should get the 2nd 2200/600 in the world, numerous SAS people were participating in this test; some CEs, John Read, and me. We worked in shifts and as foreigners we had the non-day shifts. We had a long series of tests to run and should analyze all errors – if any. But the machine ran like a dream, so there was very little to analyze.

Established in 1980

- IX4800 ClearPath test. In 1996 SAS was going to take delivery of an IX4802-5, one of the first and largest IX4800 systems ordered. Because of this I was heavily involved during the development phase and should do the final verification of the system running a SAS-setup. This ground-breaking system had many problems, so I was in constant stand-by state to fly over to start my testing. When the USAS group had successfully loaded the USAS package and run their tests, I was given the go-ahead and flew off, arriving on a Friday to use the system for test all weekend. I will never forget that weekend! I came to Roseville test floor directly from the airplane and was given initial help from Sharon Loftsgarden, who had run the USAS tests - with some problems. When I started to generate the SAS EXEC – and it failed, many times, in many ways. To make a very long story short: the system was nowhere ready for use! I got help from other testers, like Mel Peterson, special work-around EXECs, from John Carter, and even access to the latest version of the system – but I still found massive problems. I had to report this to Per Voss and others, at meetings in Eagan – and hell broke loose. Coming back to Roseville I took a lot of beating from lots of managers claiming that I had no right to be there, that it was far too early, etc. – but it all stopped when I nearly had to shout at them: we have a contract to deliver this system to SAS in two months from now! Next morning a Roseville high management meeting decided to postpone the whole IX4800 delivery program by one year -for larger dual-cabinet systems – and SAS was given a 2200/900 from Roseville as loan for a year.
- TPM testing. TIP performance monitoring is a critical requirement at high volume TIP sites, like airlines. In 1970s when USAS was introduced the standard EXEC had no good TIP monitoring, so SAS and Lufthansa [and later IBERIA] set off to define their own package, called TTL. This was later carried to the Eagan airline support team and changed a little, renamed to HVSTAT and supported in their ASP package; later again this was integrated in the standard EXEC. HVSTAT had numerous shortcomings and in late 1990s Roseville defined a much better performance monitoring package TPM. I happen to learn about this during its development and since I had lots of experience from TTL I was asked to help in the review and test of TPM. This was mostly done remotely and via mails, but I always met with John Hendrickson and Margie Smith, the main developers of TPM, whenever I was in Roseville for other reasons.

### Courses and Training

I have probably had most of the OS2200 courses, except for database and Mapper/BIS courses. I have also given more courses than I have received, so I can leave Unisys “in balance”. Some courses have been good, some not very exciting – but they have all given the opportunity to meet Unisys colleagues and most have given some nice trips

The first many courses were in London. The education center was housed in some temporary barracks in Acton – not very impressive. In fact, the walk from Acton Town tube station was very depressive as the surroundings were very badly maintained apartment blocks – or empty lots, where an apartment block should have been, now overgrown by weeds.

The U.K. education Centre later moved to Birmingham - I was never there - and then to Milton Keynes. The facility at Milton Keynes had its own hotel, so for the duration of a course you mostly stayed at the education Centre except for the once-a-week excursion to the Central Milton Keynes shopping Centre, at that time one of the largest in the world.

I have attended a few sales-oriented courses. One such was the 1100/60 product launch, where all salespeople in the Nordic were gathered in Pori in Finland. We first had to go to Stockholm, but since the company did not allow all salespeople and top management to go on one plane, some of us were put on a private plane. This was an executive jet, belonging to Muk air, and a great experience for me.





It was not a VIP plane with armchairs, but we all sat in single rows in both sides – and because I sat in the first seat, next to the supplies, I was the “stewardess” that passed drinks to the others. One other surprise was the flight pattern, where this plane climbed to 36,000 feet [now normal, then relatively high] – and then glided down to save fuel.

For some reason I was also sent on two non-OS2200 technical courses. One was a course in Vortex, the operating system on the Varian minicomputers, named V7000 in Unisys. This was held in the Varian headquarters in Walton-on-Thames, a small village south of London and it was a break from the usual Unisys courses; I never used it for anything. Another was a course in VS/9 for the 90/80 series in Brussels, where I was selected as this was considered a largescale system multi-processor system; VS/9 was a great operating system [very much like Siemens BS2000], but it was very short-lived.

In Unisys, we are always evaluated and since I seemed to be a success; someone decided that I should be selected for Unisys management training, as a possible candidate for a management role, never my desire. In the Nordics this was a series called PLU (**P**raktisk **L**eder **U**tbilding – or Practical Leader Education) and consisted of four three-day courses, all held at some nice manor house hotels around Stockholm. It was quite interesting, plus I had some nice trips – but they correctly evaluated that I was not a leader candidate!

I also managed to attend two courses in Roseville. One was a NPIT (**N**ew **P**roduct **I**nformation **T**ransfer), the other a course in Java on OS2200, the last course I received and a great event where some of the best OS2200 people from all over the world were gathered – a wonderful, but unfortunately very rare event.

As noted, I must have given more courses than I have received. When we had training facilities in Unisys Denmark, I gave a broad range of OS2200 courses, but most other courses, given in many different places, have been on TIP system operation, boot process, MASM – and XTC. A couple of these courses stand out and deserves a little more in this story of my Unisys-life.

From 1990's onwards it has been a problem to find people to support mainframe systems and specifically an “old” style environment like USAS with many programs written in Fortran. SAS at some point decided to start training people specifically for this, and Unisys Denmark provided three series of long, intensive courses, tailor-made for SAS. The first series consisted of newly educated college people, but this was no big success, as they quickly preferred to switch to some of the subjects they had been taught at colleges. Next series was people with some general IT experience, and this worked a little better, but still only few stayed with USAS. For the third and last series SAS decided to try to train some people, who were familiar with the airline world, but without IT-background; this meant that they were given a long, basic IT training before they started their Unisys training; this was the most successful.

Another memorable training session for me was in March 2012 when I was giving a MASM course – in Roseville. It was quite a challenge as an outside person to give a course within the premises where MASM was written and supported! But nobody in Roseville felt like giving the course, and as I had recently given the course, had the material and all exercises, so I was asked to give this course. Not only did I have quite a group of students in Roseville, but it was also transmitted remotely – and recorded for future use! A little later I was asked to also give the MASM course in Bangalore, India, but I believe my Unisys charge was too high, so Ron Q Smith, partly retired, gave the course instead of me - my material, I believe). Later Jan Jensen also gave MASM training in India.

Throughout many years program development and maintenance has been moved to India, but the skills transfer has been lacking. In 2017 it was decided to improve this, and Kelsey Brusio was sent to Bangalore to head up this activity. I was asked to help him out and I got a trip there to give a TIP course. It was quite an event with 60 students, cameras, dual-boards, and me wired with microphones. In 2018 I was back again to give an Advanced MASM course, then realizing what products the Indian staff had to maintain I changed the subject quite a bit underway to more useful information.

As noted under IBERIA I did quite a lot of OS2200 training for them in 2015 and 2016, gradually attended by more and more people remotely via Lync. For some courses standard material already existed [ex ECL, IPF and TIP], but for some we had to develop course material from scratch [ex. ED, COMUS, Performance], so it was quite a challenge sometimes. In my view this is really what you should do when you are on the brink of retirement: making sure younger people get trained and pick up as much as possible of your knowledge.

### TDD, UUA/E, USE, NPIT

These acronyms are a kind of technical get together, with or without customers. I have fond memories of some of these, and I miss these gatherings where you were face to face with colleagues and other people with technical interests.

**TDD** stands for Technical Discussion Days, and these were the first I attended of this kind of meetings. They were internal and brought together lots of OS2200 technical staff from all over Europe. This was in the days before mass communications, and we only had telephones (that were quite expensive for calling abroad) and telex. These meetings had two purposes: they were common training sessions, where we learnt about new features, enhancements, and new hardware. The second half was an interesting round-table discussion, where each person had the word and brought up problems and issues that they had; and very often someone would respond with the famous remark "I have a fix for that" – the "support database" of those days. The only other option we had was to call Les Shelley in London, who as the only one, had access to the PCR database in Roseville; either he would already know about the problem or he could look it up and send us a telex with the correction, so we could type it in.

**UUA/E** means Unisys User Association Europe [I think it also had other names] and was a club available for all Unisys customers in Europe. The organization was run by the customers, but it was heavily subsidized by Unisys – and Unisys provided probably 75% of the speakers. This used to be big things with 1,000s of participants in the good days, with main sessions -often with the Unisys CEO as keynote speaker and multiple parallel sessions. It also had an exhibition attached to it, where Unisys and third-party vendor provided information about their products. Often this was also the place where Unisys would announce new products. As it was a customer event it included quite a few social events also, typically with one night with an official Unisys reception, one night where each country invited their customers and attending employees out – and then some third-party vendor events. In my early years UUA/E was primarily technical, but over the years it became more and more a manager focus, maybe in line with the fact that all IT decision-making shifted away from the technocrats to the managers – and accountants! These were great events, but also hard work, as you really had to be on your toes from early morning to sometimes very late at night. With the shrinking number of Unisys customers this organization gave up and these events unfortunately do not happen anymore.

**USE** was the equivalent to UUA/E for US customers. As these events were even bigger, often much more technical and attracted a lot of specialists from Roseville then SAS preferred to send their employees to USE, they did not have to consider the airfare – and when SAS technical staff participated, I would be sent along.



In the early years these were large groups, and our salespeople went along, but in the later years I was the only Unisys person, so I had to play the role as the “host” – meaning I had to take all the bills. I must say that I was surprised how some people, when they are invited out, absolutely must have the most expensive menu item, wines, and cognacs! I believe these meetings still exist, albeit at a much smaller scale, but I do not attend.

Without UUA/E Unisys needed another way to announce new products and the concept of “Future Matters” came up. These were much smaller events and were held as smaller regional events – or directly for large customers. These often became very boring, massive PowerPoint shows, so at some point it was decided that they should have more live demos than just slides. At that point [renamed to ClearPath Road Shows] I got involved and for a couple of years, primarily thanks to Bob Hawkins, had the pleasure to develop and present some of the subjects – a great technical challenge. But they seem to have died away again – or reverted to “PowerPoint-to-death”.

**NPIT** stands for **New Product Information Transfer** and is again an internal event – and they still exist, albeit today only as web-based presentations. 10 to 20 years ago these were physical events where we met in Roseville, or maybe a European parallel session in Milton Keynes; this was great like the TDDs as it gave us an opportunity to meet colleagues and exchange ideas and experiences. They then turned into PowerPoint presentations for home study, but at least with enhanced communications the PowerPoint shows are now with a speaker – but still very boring and one-way - only written comments possible – in these ever-growing communities.

## SPECIAL PEOPLE

There are a few people that I have met in Unisys, who deserve “a place in my history” – special mentioning in my memoirs. Did I miss you in this overview? Then I am sorry, but I have met and worked with just too many great and wonderful people to cover you all.

### Tore Ask

Even though Claus Parkhøi was my formal mentor I would say that I have learned most of my OS2200 knowledge from Tore. He had worked for Unisys Norway a couple of years supporting the, then, large sites of University of Bergen, University of Trondheim, and Norsk Veritas - later Fjerndata. Tore left Unisys and worked a short period for Data General in Austria, but in December 1974 he returned to Unisys, in Denmark instead of Norway. He started at B&W, where I was, and so we got to work together. Tore had also worked in Roseville for a period, so he had lots of great contacts.

I have already mentioned, under B&W, the magic Tore performed when SECURE failed during an upgrade, where all disks were replaced. But Tore taught me many other practical details, most not documented. Especially during the Benchmarks, which we did together, his experience was of great use. We travelled quite a bit together, many times to Roseville or Eagan. Since Tore had worked there, he knew the city well, and it was also Tore, who first took me to the “holy halls of Roseville”. Tore died in 2004 – the horror year where three of my closest OS2200 colleagues died!

### Bodil Skovgaard Toftager

Bodil also died in 2004, the last of the three was Gerd Jacoby. Bodil and Gerd were rather unique in my early years in Unisys – as the OS2200 world was then very dominated by men. They were quickly “world famous” and when you travelled and said you came from Denmark it was certain that people knew about our ladies. Despite her young age Bodil and Gerd too had quite an experience, both having worked as very young at SAS. Bodil had also worked at Air Canada and that is why she was an outstanding expert in TIP. Before she came to Unisys Denmark she worked at SHAPE in Belgium, so she had quite a C. V.

I met Bodil when I came to FLS where we had to install TIP. Taking advantage of Bodil's experience FLS was selected as one of the sites to Beta test EXEC 33 where TIP was integrated into the EXEC - before that it was a separate package developed at Air Canada, where TIP files were independent of EXEC files, residing on its own removable packs. This was quite a challenge and Bodil and I spent many hours, often until late at night, to test out the new TIP and debug it. So, we came to know each other very well. We also used Bodil's TIP knowledge for the Tjæreborg benchmark and Bodil joined me and Tore in Eagan.

#### Michael Deneyer

When Michael joined Unisys, he was sent to FLS to work with me, so I became his mentor. He never forgot that and always praised me, much more than I probably deserved. We also had great time together socially – until Michael also died from cancer, far too young.

#### Anders Öhman

Another 'pupil' who became a close friend was Anders, who came to work with me at SAS in 1990s. Anders was an experienced OS2200 person, having supported the OS2200 systems for the TELCON development Centre in London. Before that he had worked at many sites in his native Sweden, mostly working on the OS2200 communication products. So, the airline environment was new to him. We worked together for two years, and he was without doubt the best person I worked with during my years at SAS. The SAS staff also liked him a lot, and in SAS - where many Danes in SAS disliked Swedes as they felt they ran SAS only their way - he was the unique Swede they really liked – and disturbed the picture of "the bad Swedes".

When SAS cut down on Unisys staff, Anders worked a few years at other Danish sites before he moved to Spain and took over responsibility for IBERIA. Anders was very liked there too and ran the site with fantastic capacity and made the site a flagship Unisys OS2200 site. Anders ran IBERIA very much along the same line I ran SAS, so for many years we shared lots of experiences and took advantage of what each other did. When Anders retired, I was often asked to come and help at IBERIA, which gave me many nice trips there.

#### Palle Hansen

Palle was for many years the EXEC person and group leader for the Unisys support staff at DSB. Like the group at SAS, they lived their own life at the customer site, so it took a few years before we worked together – except when we suddenly had to help each other in emergency situations. Palle later moved into the management ranks [without ever forgetting his OS2200 roots] and Palle became my de-facto manager for years. He still is – and the reason I decided to write this document – without actually being so, due to the complicated Unisys organization structures.

Now, after being laid off from Unisys and moved to Miracle, Palle has become my direct line manager.

#### Jørn John Loberg

As noted at the start I probably should have worked in Jørn's group in Unisys but was captured by the OS2200 manager. We met frequently in Unisys in the early years, and at some point Jørn moved to work for F.L.Smith, where we also met work-wise. However, it was mostly in England we got a lot together with Jørn and his wife Manja. Jørn worked in the Telephone Assistance Group in Bridge House, same as me, and he also had a house in Maidenhead, not far from us.

Jørn later moved to Roseville where he continued to work, despite being 70+, in the EXEC Continuation Group and solved lots of customer-reported problems. Jørn was laid off from Unisys in the cut in 2015 – at the age of 74!



### Joe Nolan

I met Joe in Bridge House, where he was heading the Quality Assurance group, responsible for testing the AIDS (later COMUS) products. This was a very careful and detailed testing; I am sure this is one reason why COMUS has been so reliable and is still around 40 years later.

In 1982, when COMUS was to replace AIDS, it was decided that Joe and I should oversee the conversion of all COMUS1 and AIDS product build procedures to the new common COMUS2 version, 28 products in all. This project took place in Roseville and was scheduled to last five to six months, so both families moved to Twin Cities. The project work is described above, as well as some of the social aspects on being in the U.S.

After the COMUS project was completed, we kept good contact. Joe returned to Roseville and worked on new systems and later he oversaw the system support group at our giant installation at UBS in Zurich – commuting back and forth every week, had his own seat on the plane! Joe retired some years ago – after the closure of the UBS project.

### Ian Hepburn

I first met Ian at RECKU and later he became my supervisor in the STI group in Bridge House in London during the AIDS project.

After the conversion of AIDS to COMUS Ian stayed in Roseville and worked on numerous projects. To me he has been the “inventor” of two very significant features in OS2200. Shortly after moving to Roseville Ian asked me to review a document that he had written, where he suggested how to update the OS2200 library files on-the-fly – before this they had to be loaded on the boot tape and the system had to be booted. This was the start of AFCBs, the ability to update nearly all software products without stopping the system – something Microsoft is not able to, 30 years later. The other great idea that Ian fostered was console automation, which over the years grew into the SPO (**S**ingle **P**oint **O**peration – or Operation Sentinel).

Ian later moved to Seattle and worked in the close cooperation Unisys has had for years with Microsoft [I think, I never knew very much of that part of Unisys business]. Ian has now retired, but still lives on the West Coast, partly in Seattle, partly in Long Beach.

### Rick Onsgard

Apart from Jørn, with his Danish background, Rick is my closest colleague and “ally” in Roseville. Today Rick is “Mr. TIP” in EXEC continuation, so we often work on the same issues and problems. Rick is very helpful and a considerable part of the success I have had at largescale transaction sites.

There have been several instances, where Rick has been the one to help me get through troubles at TravelSky. Sitting in the TravelSky machine room I have more than once processed dumps on my PC and sent selected bits and pieces across to Rick and we have been able to overcome, or even fix the problem that way, long before we could FTP a dump file across. If Rick has not been the primary person, then he has often been the remote coordinator.

If Rick is so valuable for TravelSky then why is he not onsite for critical upgrades? Well – he has been once. But Rick is a big guy, so flying 11 time zones is very uncomfortable in an economy class airline seat. After some talks, TravelSky and/or Unisys China accepted that Rick flew Business Class to Beijing, and Rick provided excellent support. But after the trip Rick had massive problems with all kinds of Unisys managers and other approvers to get his expenses for the trip approved – so Rick says: Never again! This is how stringent rule, and unflexible managers make sure that we cannot provide the best possible support for one of our most important customers!

I can always call Rick for help. I owe him a lot.



Rob Jamieson

I have met Rob in connection with TravelSky. Rob comes from Brisbane in Australia, like several other people I have met along the way and has worked at Lufthansa. In fact, Rob and I had, a few years apart, the same room in the house in Schwanheim by Frankfurt.

Rob is fantastic at the parts of OS2200 software that I know very little of, such as DEPCON (EOM), Eclipse, ePortal – and he is outstanding with SPO, where he has developed numerous valuable features, mostly for TravelSky, together with Jim Malnati. Rob was a very large man, with some heart problems. But a few years ago, he started running, lost 50kg and is now running marathons! He runs a lot more than I have ever done.

## EPILOGUE

*I had a great time with Unisys.*

I think Unisys is a good company, but most of all for me Unisys is all the wonderful people I have met and worked with! I have met a few Unisys people I did not like, [mostly people who talk as if they have invented everything in Unisys on their own], but then I just avoid them. Otherwise, I think Unisys somehow has been able to employ great people. From family and friends, I hear about different work environments and different workplace “tones”, so I appreciate the managers and colleagues I have had. Yes, there is teasing and fun in the workplace, but it has always stayed at a very sober level, from my experience.

Was everything wonderful? No – of course not. Work never brought me down, but management reporting has nearly done! I have been on the brink of leaving Unisys several of times, simply because of the yearly (or nowadays twice a year) filling out of stupid reports about goals, development, experience and whatever! My good and very clever colleague Jan Lauridsen left Unisys partly because of this stupidity – we lost a good man there!

Web time also has me upset lately. Web time itself is OK; except I find it hard to accept a system that requires hours of outage every week – if we delivered this kind of system to our customers, we would have been out of business years ago. However, the requirement to report on a Friday afternoon is absolute stupid – at best. As a permanent customer support, I am, in principle, always at work and can be called anytime – including Friday night and all of Saturday. The reporting period runs to midnight between Saturday and Sunday and reporting before that time is guessing at best, lying to say it correctly. How can that work together with our permanent yearly goal to report time accurately?

Did I miss anything while working for Unisys? Only a few things. Over the years, I have heard stories of these wonderful Club events for the salespeople that have met their quotas – and that non-salespeople were sometimes invited; but I never was. There has been countless of instances where we have saved a sales representative who has promised “the impossible”, and in the end we made it work – and they therefore made it to the Club meetings!

Two other misses: 1) I never visited St. Paul de Venice! I was supposed to attend numerous times, but I was always dropped off the list at some point. And finally: 2) Unisys has sent me to so many places, but never to Australia!

*But in the end: Unisys – you have been very, very good to me!*

If You are mentioned herein, email [la.gj.benson@comcast.net](mailto:la.gj.benson@comcast.net) for the original paper copy.

