

Offering large-system machine translation as a bureau service

Peter Wheeler

Antler Translation Services
Sparta, NJ 07871, USA

This is an essentially personal account of the work of my company in providing large-system machine translation as a bureau service. It does not describe a research project, carefully funded from the outset and following well-defined investigative guidelines. Still less does it set out to be a blueprint of how such a bureau service should be run. It is a record, warts and all, of what actually happened, in the hurly-burly of the marketplace.

It is certainly the case that if it had been a research project, things would have needed to be planned much more rigorously than they were, and similarly in a theoretical description of how such a service should be run, some of the decisions would have been taken differently.

But that isn't the way it happened.

First, some background. My machine translation career started, like many other people's, right here in Luxembourg, shortly after the Commission had bought certain development rights to Systran.

The initial English to French project, and the whole overall effort, was of course led by Ian Pigott, to whose dedication and dogged - not to say obsessive! - determination to make it work many of us here owe a great deal. In late 1976, while I was working as a translator for the Commission, he was looking for someone to take over the embryonic French-English system. Like all other professional translators of the time, I "knew" that machine translation couldn't work - we had all been raised knowing that there was something inherently intractable to computer processing about language - but it suddenly occurred to me, in one of those flashes of thought that come we know not whence but which are to have such far-reaching consequences - that I ought to hear what Ian had to say, and then I would "really know" that machine translation couldn't work, and I could forget about it along with all of the other impossible dreams.

In that spirit of scepticism, then, I went along to the interview..... and was instantly captivated. Not only did I see that machine translation might work after all, I could also see that it would be an inordinate amount of fun to try to make it work.

From the vantage point of almost 20 years later, I see us as being then like the amateur explorers of the last century: we had no clear idea of where we were going, our equipment by modern standards was laughable, but in the boundless enthusiasm of the British amateur we thought it would be jolly good fun to have a crack at finding the source of the Nile or taking a run at Kilimanjaro over the weekend.

In those early days, Ian and I didn't even have access to a screen and a keyboard! We would write out our coding instructions on coding sheets, they would be punched to tape by a small team of typists, and the loading of the instructions and the subsequent running of the test translations would take place on the Commission's mainframe, under the control of a programmer who would select the requisite IBM punched cards largely - it seemed - at random and from memory, often with disastrous results.

The translations and updates would come out on huge printouts, which were soon piled haphazardly all around our office. It seems almost primitive to think of it now, but it must also be remembered that the translation department - and I still considered myself as very much a translator, working to develop a useful tool for my colleagues - was still in the typewriter era also. Those translators who typed had Commission-issue manual machines, the typing pool had electrics, but personal computers and word-processors were unheard of.

In fact, - and this is something I remembered only when I began to write this paper - the first Wang word-processors in the translation service's typing pool were put there precisely to support the Systran translation effort. We ultimately got some word-processors into our DG XIII development effort, which were hooked up so that we could use them as remote job-entry terminals and start running Systran jobs ourselves, and later a couple of them, and then more, were seconded to the English and then the French typing pools. I became the Wang guru, not only teaching innumerable secretaries and translators how to use them but also physically moving machines around and laying cables from one office to another.

Nowadays, when screen-based machines are commonplace and the Commission even has a demonstration room, with comfortable chairs and professional displays on which to show off its various computer applications, it seems incredible to remember my doing a Systran demonstration to a group of students crammed into my small office, with my Wang screen perched precariously on a chair and the chair itself standing awkwardly on my desk!

Our contacts with our various subcontractors were equally ad hoc. In those early days we had a group of freelance dictionary coders working directly for us, so directly that Ian would personally count the number of entries they had made, subtract the ones which he did not consider valid, and pay off the contractors at so many francs per entry!

Gradually, over the years, things became less informal. A formal Systran development body came into being, and existed in a variety of locations, under a variety of leaders. It worked for us under a formal development contract, but the linguists we were actually dealing with from day to day were - surprise, surprise ! - the same ones as before. The machinery became more sophisticated, the development effort became more organized, and above all the systems did actually go into production. Some of the English and French translators (and later other languages as well) became enthusiastic at trying out this new technology.

As some of the amateurishness went out of the project, however, some of its fun went too. And so I was ready for a change when Systran's big rival Logos asked me to come on board. I had always had an unreasoning fascination with America where, I thought, all the women would look like Marilyn Monroe and Faye Dunaway and I would instantly turn into a combination of John Wayne and Humphrey Bogart, and so I agreed.

America was not to be for another year, however, and so for 12 months I lived and worked in Frankfurt, commuting back here to Luxembourg on Friday evenings.

Frankfurt housed the headquarters of Logos' European operation, and my role there was to handle dictionary development for German to English. In that role I had the help of some staff, but they were 3000 miles away in Middletown, NY and Boston!

Logos at the time was a series of autonomous communities, which never quite gelled together. In the order in which I encountered them, there was firstly corporate management in Boston - the dark suits, the total concern with the stock price and finding the money to keep the company afloat, the incomprehensible jargon (when the then CEO said to me that "Systran has an RFP out on the street, doesn't it?" I just grunted noncommittally, not knowing

what an RFP was or why it would be out on the street!) Then there was Logos Frankfurt: mostly young, mostly German but with a smattering of Brits and one American, all speaking several languages, all very dynamic and very committed to selling this new technology to German industry and seeing it work, and all ultimately frustrated by what they saw as a lack of responsiveness from the development centre in Middletown.

Which is where I ended up moving to (not Boston, as had originally been planned). Where I found the mirror image of the frustration, only this time it was at Frankfurt's impatience and inability to understand the complexities of the development process. But at the same time I also found that Middletown was itself an uneasy mixture of two different communities, one of which had been quite literally a community or commune, living together under the spiritual guidance of two men one of whom had had the idea of the Logos system. And although what they had achieved was remarkable, they shared a mindset which delimited them very sharply from, and excluded, the outsider community which at the time I moved to Middletown consisted only of four or five people. Outsiders were people who knew several languages, knew more about machine translation than just Logos, people who had been to other countries.

(They were also incredibly secretive - when I mentioned that before I came to Middletown I already knew that their research centre was in an abandoned railway station, they were aghast that someone had talked!)

During the course of my years with Logos, the composition of the company changed. There was a great influx of people with a more conventional background in linguistics, or programming, or documentation production, so that while the numbers of the original commune-members declined only slightly, their proportion and their influence dropped sharply.

Meanwhile my responsibilities grew. From German-English dictionary development, I then took over all responsibility for German-source systems and later all English-source systems as well. By this time, my title was Group Product Manger, with overall responsibility from a linguistic and marketing point of view - for ensuring that they were of use to the translators who would be using them - for all products leaving the development centre. I was, in short, the second in command to the Development Director, and deputized for him while he was away.

Meanwhile, in the image which it presented to the outside world, and the way that its products were marketed, Logos became considerably more professional. We had installations in six countries (including, briefly, the Commission here in Luxembourg). We moved from the railway station to a proper office building, with real air conditioning (the company sold off all of the fans it issued us in the summer and the heaters we had for the winter, and one of them to this day is used to warm up one of my cars in the winter!) and a real computer room with a proper raised floor. We expanded within that building, and were a week away from expanding a second time when the axe fell.

In one dreadful day, about three-quarters of the linguistic staff were given notice to leave.

The money had finally run out. You will forgive me if I do not talk in detail about this painful time.

There remained the question of what to do with the Group Product Manager, who had seemed a reasonable expense when he was managing a staff of 46, but who was now much too costly when his team was about ten. The first attempt to move me out actually gave me

a chuckle. Logos: "We'd like you to leave on June 30, and before you go would you please assign someone to go and install a system in Kansas City in the first two weeks of July. Someone who knows the Wang VS, who knows how to give training, and who knows both French and German." Reply: "We are now so thin on the ground that there is only one person left who meets those four requirements: me!" Which meant that I stayed on for another two weeks and had a quite pleasant trip to Kansas City into the bargain.

Finally, however, I was out of there. But at the same time not out of there, because what I had decided to do was set up a company which would market Logos translations to clients not large enough, or not yet ready or willing, to install a system of their own. At the time, the idea was unique. Systran in La Jolla had (and still has) an arrangement whereby users can dial into the Systran computer and run their own translations, but my idea was different: the client with very large translation needs would turn the jobs over to me, I would handle all of the dictionary processing and other preparatory work, run the translations on my own Logos system, and return the finished product to the client.

Nor was I physically out of there: in a very fair severance arrangement, Logos were allowing me to operate within their office building (where now, of course, they had far too much space!) The importance of that to getting my company off the ground was immense: I had no money (former colleagues here in Luxembourg will recall that I never had any money even when I was earning the same enormous Eurocrat salary as they were) and yet through my deal with Logos I was able to have my choice of office furniture (they had too much of that too), plenty of computer power, the use of a fax and a coffee-machine, even a receptionist to take my phone calls!

I even had a Logos system, running on a mainframe! Well, I said I did, anyway. What I touted to my potential clients as "my" Logos system was actually Logos's own development system, running on their computer. If and when I actually got any business, then I would pay Logos an agreed fee, so many cents per word, for the processing on their machine.

Thus the arrangement was actually beneficial to both sides. In return for their making available free office space - which they had to pay for anyway - Logos had a chance to make some money back on some of their idle computer capacity. Similarly, as I was prospecting for clients if I found a company which looked like a candidate actually to buy a Logos system, then obviously I would tell the Logos sales people. Likewise, if their salesmen found someone whose need was too small to justify the purchase of a system, but who still had large translations to be processed, then Logos would pass them on to me.

Which is, indeed, where my first client came from. Trying to sell a Logos system to one of the local branches of the huge Price Waterhouse accounting firm, Logos found that PW did not want to install a system, but that they would be interested in having me run a translation of one of their management software manuals.

This was the job that proved that in outline, at least, the idea was going to work. There **was** a market out there. I took the client out to lunch, pretending I knew what I was doing, I negotiated, I got the job.

Flushed with that success, I embarked on a massive expansion of personnel, doubling the size of the company overnight! In other words, I hired one more person! Logos' former head of Spanish lexicography joined me to handle sales and marketing, and is with me still.

We ran the Price Waterhouse text through New Word Search and Noun Phrase Search, we added hundreds of management terms to the English-Spanish dictionary (terms which then

became Logos' property, as part of our deal), and we ran the translation on "our" Logos system. Then the raw translation went out to freelance post-editors.

It was about this time that Logos was changing its front end from the Wang word processor to the PC, and so the output coming off the mainframe was converted to WordPerfect files and sent to professional translators who had WordPerfect (which, in the US, virtually all of them do). They post-edited on their own computers, the work was shipped back to us on diskettes, final quality control was run and some more dictionary work done for the future, and the completed job was shipped off to a satisfied client.

The fee, I recall, was just under \$20,000, which seemed to me at the time a **vast** amount - I was on my way to fame and fortune!

I wasn't, of course. While the Price Waterhouse job was being processed, we were pursuing other leads, most of whom teased but then said "no." I remember doing a huge mailing (with the addresses being typed up by another ex-Logos-colleague, now unemployed). Response - absolute zero.

It was a very typical startup time for a small company - every little success seemed like an absolute triumph, the key to glory, every little setback was a sheer disaster. A time of scratching around to make ends meet - scratching around in my case taking the form of working for Dragonfly Software and for the UN - a time of wondering where the next hundred dollars was coming from.

Where the next hundred dollars was coming from, again, was another Logos no-thank-you. A large chemicals company with a database of warning statements. Which was to be translated, initially, into French.

This was the ideal machine translation application. Not only was it already available in machine readable form (as the Price Waterhouse material had been also, of course), not only was the language relatively simple ("Do not mix this chemical with water;" "This material explodes on impact;" "This product is not suitable for human consumption") it was also *extremely* boring:

Causes vomiting.

Causes vomiting.

Causes vomiting.

Causes headaches and vomiting.

Causes headaches and vomiting.

Causes nosebleeds and headaches.

Causes nosebleeds and vomiting.

... and so on, for hundreds of pages! Not the sort of thing a professional translator particularly wants to spend their time doing, and the sort of thing that is mind-numbingly tedious to concentrate on for hours at a stretch. But provide the translator with a first draft, written by a not very bright but extremely fast machine assistant, and the work just zips along! The procedure was as before: New Word Search, hundreds and hundreds of chemicals, poisons and diseases added to the Logos dictionary, raw translation and post-editing by freelance translators.

This was the job which refined the concept, taught us more about what we were doing. Reinforced the lesson, in particular, that even when the machine is doing a lot of the work, there is no room for mediocrity in the post-editing stage. Because:

Causes headaches and vomiting.
Causes nosebleeds and headaches.
Causes nosebleeds and vomiting.
Causes nosebleeds and buffalo hump.

What's buffalo hump? Since the phrase had not been specifically entered into the dictionary, Logos translated it literally on the basis of its component parts: *bosse de buffle*. Sounds convincing, and since much of this terminology is literally identical from one language to another, a careless post-editor might have let it go. However, the post-editor we had on the job, not only a professional translator but a licensed French pharmacist, knew that the correct French for the disease "buffalo hump" is not *bosse de buffle* but *bosse de bison* - "bison hump" - and a disaster was averted.

Another satisfied customer, another dent in the overdraft! Other smaller dents were made by other clients, one in particular a translation agency in New York which was interested in trying out machine translation for some of its jobs. To this day a major Japanese auto-maker does not know that one of its handbooks was partly translated using Logos!

The other memorable thing about that client was that he wanted us to transfer files backwards and forwards by modem. Again, we have to remember how long ago this was - most translators did not yet have modems, I didn't, surprisingly even Logos didn't. But as the head of a cutting-edge company selling the very latest thing in translation technology, obviously I couldn't admit that I didn't so I invented the story of the employee who had tripped over the cable of the PC which contained the modem, sending it crashing to the floor and putting the modem out of commission!

Later, the chemicals company came back, to have the same database translated into German. This time we could make use of much of our prior work, since the source side of our English-German dictionary entries existed already from our English-French work, and at the same time we knew at what points in the text the analysis component of Logos was likely to have difficulties.

It was about this time that the situation took an unexpected turn.

As part of the ongoing effort to diversify, to find some income somewhere, anywhere, I had become a dealer for the Globalink system, and was thus in the position, which may at first seem strange, of selling Globalink machine translation packages from within the Logos machine translation development premises! But I thought then, and I still think, that there was nothing unethical about my position: I was not an employee of Logos, I was a tenant, and indeed in some situations we had actually been competitors. And above all, Logos and Globalink are simply not in the same market: the buyer of a \$500 Globalink package is just not a potential client for a \$20,000 Logos system.

Unfortunately, some of the Logos hierarchy did not share my philosophy, and when they found out about my marketing ploy they threw my company out of the premises! An episode which encapsulated the ambiguity of our relationship: the Logos employee on the one hand hounding us to return our keys for the building, and on the other remarking quite conversationally that it was possible to open the door with a plastic credit card; the Logos comptroller writing a really stern letter giving me less than 24 hours to get out of the building, and then cheerfully negotiating the sale to me of some furniture and other supplies; and above all the Logos salesman calling up a month or so later with the opportunity we had all been waiting for.

General Motors was looking for a machine translation company to translate its data base, about a quarter of a million words, into French for the Canadian market. Were we interested?

Negotiations followed, we won the bid, we set about translating in just the same way as before. Once again, an ideal machine translation application, in that, once again, it was extremely tedious. A series of disjointed statements, making little coherent sense on their own but joined together as needed by the GM computer in order to put together a manual.

Like "Your car has <codes> four <codes> five <codes> six <codes> wheelnuts." "The petrol tank has a capacity of <codes> fourteen <codes> fifteen <codes> sixteen <codes> gallons."

But while the content itself was sometimes a fraction on the dry side, the database layout, with its choice of options, raised a whole new set of problems. What to do with a case like "Press the <codes> clutch <codes> clutch pedal <codes>," for example? Since in French "clutch" is masculine but "pedal" is feminine, something creative has to be done with the position of the article.

Or "This switch is located on the <codes> left <codes> right <codes> door." Once these codes had been made invisible to Logos, then the sentence would translate correctly, but when they were put back into the text there was a need to check that they had gone back in the right place, so that they were offering an option between *gauche* and *droite* and not, say, between *gauche* and *porte*. At the present time, some years later, Logos has become much more sophisticated at handling markup codes such as SGML, but back then we largely handled them ourselves by some wordprocessing routines and a lot - an impossible amount - of eyeballing. Many of these problems, of which these are just two simple examples, had not occurred to GM before they decided to undertake the translation, and they learned about languages as we learned more about databases!

It was during this project that I learned forcefully that even with the machine's help, these big projects can sometimes be almost overwhelming: one particular piece of the puzzle, having more to do with the database codes than with language, took two days and two nights of continuous work, after which I had to drive through the third night to a conference, just like this one, to talk about the joys of machine translation!

By now, the company was in a paradoxically ambiguous situation: apparently, a slick and successful supplier to the great GM (the cachet of whose name was to open many doors in the future) which on the telephone, to the inquirer who could not see us, sounded extremely professional, like a real company. In reality, two people working out of a tiny study in my house, with enough desk space for one keyboard (the two PCs piled one on top of the other, and the second person having to work with their keyboard on their knees); the cat playing around in the mountains of paper about GM cars; my daughter's horse-riding gear piled up on the wardrobe-sized controller which is what connects us to the Logos computer; and at times the neighbour's teenagers playing their music so loud that we can't hear ourselves think.

But we got the job done.

We delivered the database on time and on budget.

The most visible consequence was that I could afford to move the company back out of the study in my house, and into a proper office. Which was to prove serendipitous when our contact at GM called again in the following spring: "We're going to have some more work for

you, but my bosses want me to come and look at you. They want to be sure that you're a real company, working out of real premises. Of course, personally, I don't care if you're working out of the study in your house " We laughed man-to-man, but only one of us knew that that was indeed what we had been doing all the previous year!

So the great GM was going to come and see us, watch us hook up to the Logos computer, watch us do dictionary updates and run translations. It was all going to be very impressive, but suppose they thought that this two-employees-and-a-temp company was too small to handle the rest of the business? (Which we had since learned was going to be 28 owners' manuals, again for Canada). What were we going to do?

We did what small companies do, we invented phantom employees. An unoccupied desk with a photo of my daughter on it, some papers scattered around, a bottle of nail varnish left carelessly on the desk, and we have an image of another employee who has just painted her nails and gone out to lunch. An extra man's coat hanging on the coat rack, and there's another one.

Whether the visitor was fooled or not, GM gave us the business! But already a change was setting in in the way we processed the work. We began to realise that quite apart from the pieces which were in the database, and therefore common to many manuals and translated only once, there were also many other parts which were the same, or nearly the same, from one manual to another, but which we were translating over and over again. Gradually, we developed a system for recognizing such similar parts, and, once they had been translated in one manual, re-using the translation, perhaps with minor edits, in another. We were saving time, we were saving money.

And we were foreshadowing a later trend in the whole process of computer assistance to translation, namely the development of translation manager systems to recognize and pass through already translated sections.

In our case, though, it has to be admitted that the "translation manager" which we were using was largely my own memory, based on the fact that I had personally read every word which we translated for GM at least once, usually twice.

Subsequently, however, we began to realize that even where the bulk of a section was unique, there were phrases in it such as "Turn on the headlamps," or "Open the windows," which came again and again. So why not pre-translate them? With a number of basic search and replace routines, incorporated in a simple program which would call them one after the other, we began this rudimentary pre-processing. The editors, then, who were now something of a hybrid between machine-translation post-editors and translators, might receive a sentence which read "If it becomes dark, *allumez vos phares*" or "If the windows begin to steam up, *baissez les vitres*"

Then we went further. We might find, for example, that the manuals always rendered items in a different order between the English and the French. "If it is raining, turn on your headlamps," for example might always be rendered not by *S'il pleut allumez vos phares* but by *Allumez vos phares s'il pleut*. If it was not possible to handle the whole sentence in one go (remember, we were working only in wordprocessing, with some limitations), then we would create interlocking change routines whose individual parts looked like nonsense but which made sense as a whole.

We might make our system translate "If it is raining" by *Allumez vos phares* and "turn on your headlamps" by *s'il pleut*, although of course we had to ensure that such weird translations were called only where both halves occurred together.

And the way we would ensure that these translations were picked only when required was equally pragmatic: we would actually write expressions like "If it is raining turn" to be translated by *Allumez vos phares* and "on your headlamps" by *s'il pleut*.

Of course, there are dangers with such an approach. If the manual said somewhere else, for example, "The correct bulb wattage is stamped on your headlamps," then the editor would receive a sentence that said "The correct bulb wattage is stamped *s'il pleut!*" Which did have the side benefit of keeping the editors alert!

I am sure that such an approach will make real computational linguists cringe and shudder, but as I said right at the beginning my aim was not to build a machine translation system, not to lay down rules for how a machine translation bureau service should operate, but just to keep my business going out in the real world. And as we built these pragmatic tools we were not only foreshadowing the development of translation memory tools, which is where the interest of the professional translator community now largely lies, rather than in more "pure" machine translation systems, but we were also discovering a deep truth: that out there in the real world, the customer really doesn't care how the job is done.

The customer wants his translation in a given format, by a given date and at a given price. How the text gets to that state is largely a matter of indifference to him. With hindsight, I came to see then that part of the reason that finding the clients had been so difficult in the first place - apart from those who actually came looking for us - was that I had scared them away with too much high-tech. For many companies - although this is a remark from the perspective of the US, rather than of Europe - it is enough of a new idea that they have to translate their publicity and technical material into a foreign language at all, they don't want also to be dealing with ground-breaking questions of machine translation, semantic processing, subject-specific dictionaries and the like.

And with that insight, things seem to be working out fine: whether we translate by genuine machine translation, by some word-processing hybrid, or by using real-live human translators, the clients keep coming back!