David Gugerli | Ricky Wichum

Simulation for All
The Politics of Supercomputing in Stuttgart

CHRONOS

zurück

David Gugerli, Ricky Wichum

SIMULATION FOR ALL

THE POLITICS OF SUPERCOMPUTING IN STUTTGART

Translator: Giselle Weiss

CHRONOS

zurück

Cover image: Installing the Cray-2 in the computing center
on 7 October 1983 (Polaroid UASt).
Cover design: Thea Sautter, Zürich

ISBN 978-3-0340-1621-6
E-Book (PDF): DOI 10.33057/chronos.1621
German edition: ISBN 978-3-0340-1620-9

zurück

Contents

zurück

User's guide

The history of supercomputing in Stuttgart is fascinating. It is also complex. Relating it necessarily entails finding one's own way and deciding meaningful turning points. Many accounts thus begin with an origin story whose impact is said to persist to the present day. Origin stories satisfy the desire for maximum simplification. The history of supercomputing in Stuttgart is so complex that it has given rise to two different origin stories.

One of these stories involves the appointment of aerospace engineer John Argyris to the University of Stuttgart in 1959.[1] Beginning with Argyris means tracing contemporary simulation culture as far back as it will go. Argyris wished to apply numerical methods to problems of elasticity theory. For this purpose, he used a mainframe Ferranti computer evocatively named Pegasus to test different solutions. This origin story gives rise almost effortlessly to a narrative that portrays state-of-the-art computing in Stuttgart as an interplay of local advances in methods and deployment of the fastest, most powerful machines.

The other, no less mythical origin story of Stuttgart supercomputing is the "pivot to the future," which Lothar Späth – ministerpresident of Baden-Württemberg from 1978 to 1991 – hoped to usher in by purchasing a Cray-2.[2] Installed in Stuttgart in 1986, this machine was the fastest computer in the world and the first of its kind in Europe. But the Stuttgart Cray-2 was also the product of an unprecedented procurement process that sidestepped the complex negotiations characteristic of funding agencies and relegated the state and federal governments' coordinating bodies to idle spectators. Späth is the referent for narratives in which a strong minister-president and a bold university rector are vested with decision-making power in supercomputing matters.

Our study takes a history of technology approach to the story of supercomputing in "Stuttgart", i.e. the university, its computing center, the academic and industrial region or the government of Baden-Württemberg. We focus on the problems to be solved and the solutions that have enabled Stuttgart to perform at the limit of computability for half a century. We take the position that origin stories are ill suited to recounting this history. Its course of development was too complex and punctuated by too many interruptions, crises, "reboots," and surprises. We believe that supercomputing in Stuttgart was reinvented at frequent intervals to maintain its appeal as a service to science and industry. The fundamental building blocks of this service included operations, user policy, financing, and science policy. These building blocks were continually shuffled and reshuffled at the computing center. A computing center itself is the product of the interplay of machines, networks, buildings, personnel, and users; of competition and cooperation with other computing centers; of the interests of the sciences; and of the shaping power of funding agencies, university administrations, and industry.

How did the vicissitudes of science policy influence supercomputing in Stuttgart? How did the switch from vector to multiprocessor systems change Stuttgart's simulation culture? Which sciences influenced the development of supercomputing? What role did industry play in Stuttgart? How were users introduced to supercomputers, made familiar with them, and simultaneously tutored in their use? These are the questions we wish to address. In so doing, we want to highlight problems occasioned by conflicts between local operations and trends in science policy that

Fig. 1: The work of reconfiguration: rarely seen, but essential.

affected supercomputing in Germany and elsewhere in Europe, and to show how those problems were solved.

We have divided our history of supercomputing in Stuttgart since the 1970s into four sections. In each section, we have identified specific organizational, scientific, and technological strategies that brought computers, personnel and programs, buildings, networks, and users, as well as the institutional and political framework, into a new relationship. As elsewhere, the history of supercomputing in Stuttgart constitutes a series of configurations.[3]

In the first configuration, much of the discussion centered on the highly controversial centrality of the service. In 1972, the University of Stuttgart Computing Center became the Regional Computing Center of the University of Stuttgart (RUS). The fastest machine on-site was a CD 6600 made by Control Data. The mission of RUS was to supply the regional universities with computing capacity. To this end, RUS lobbied the Deutsche Forschungsgemeinschaft (DFG), Germany's national research foundation, to extend its building and computing capacity. Nevertheless, in 1983, more by accident than by design, a Cray-1 – an aging supercomputer from the 1970s – was installed at the computing center.

The second configuration emphasized local supercomputing performance. In 1986, the first Cray-2 on the European continent landed in Stuttgart, with great fanfare and through political force of will. It was a spectacular acquisition. Yet the computer was difficult to operate, and procuring a replacement turned into a project lasting several years. Supercomputing was increasingly about performance, testifying to West Germany's global competitiveness. Stuttgart's procurement problems remained unsolved until 1996, when the High Performance Computing Center of the University of Stuttgart (HLRS) was founded and commissioned to provide peak computing capacity for the entire federal republic. The HLRS was financed by a public-private partnership that also involved the regional energy and automotive industries. The HLRS was organized as a service center responsible for distributing computing capacity.

The third configuration depended very strongly on the combination of heterogeneous computer architectures. By the end of the 1990s, the HLRS was not only coupling vector computers to massively parallel computers: it was even planning to offer meta-computing to the entire world. At any rate, establishing connections between institutions and local operations opened the way to the integration of German supercomputing into a national research network (D-GRID), though it would take a considerable amount of cooperation, conceptualizing, and committee work.

Much of the work in the fourth configuration focused on supercomputing users. Training programs were attracting more users than ever to the center, and service-level agreements defined a new relationship between users and the center. In 2006, the HLRS moved to a permanent address in the new building at Nobelstrasse 19 on the Vaihingen campus. Owing to its expertise in virtual reality simulation, the HLRS became a cornerstone of the University of Stuttgart's structural realignment. At the same time, the university was striving to shift the spiral of innovation and investment in high-performance computing to the European level.

zurück

The centrality issue (1972–1987)

Gaining dominance

Computing centers must go to great lengths to secure their claim to supremacy. That is why university documents are given over to so much discussion whenever a computing center has been newly established (or expanded, relocated, or otherwise organized). Historically, this has been as true for the "data factory" type of computing center as for centers specializing in peak-performance or "high-performance computing." None of these facilities could count on their "centrality" being guaranteed over the long term, and all of them had to contend with a heterogeneous environment.

By the early 1970s, the University of Stuttgart had established a fair degree of centrality. At the time, computers in various institutes were being dismantled and replaced by one large computer at a central location. This move was accompanied by adaptations and a certain amount of replanning. Established locations were downgraded to outposts with a fixed expiry date. Core pieces of new equipment were given a prominent address. Personnel were redeployed. Existing connections were severed, workarounds introduced, and shortcuts proposed. In addition, the organizational chart was redesigned, and responsibilities were redefined.

Such is the genesis of a computing center. The same pattern of actions has been repeated time and again, sometimes with a view to the future, sometimes in retrospect. For example, a 1971 University of Stuttgart report on the use of computer systems states that in 1958 the Technische Hochschule had already purchased the

zurück

first German digital computer, a ZUSE Z22, and simultaneously "founded the university's own computing center."[4] In 1968, the realization "that running one large machine is more economical than several small ones working together" led to the founding of the "Regional Computing Center," where "a CD 6600 was operated jointly with Control Data [Corporation, CDC]."[5] The center had apparently been founded quite some time ago and been active for a long while.

In recapping events and computing centers of the past, the university's 1971 report was presumably anticipating further centralization and creation. Indeed, the next year's annual report states that 1972 was "a year of organizational new beginnings" for the university's computing center: "The former University Computing Center and the Regional Computing Center were merged … to form the Computing Center of the University of Stuttgart (RUS)."[6]

The 1972 merger did not introduce a simplification. Rather, it involved strategic separations, such as into two departments. Department A was now located in the city center, and Department B in Stuttgart-Vaihingen. At the same time, the university parted company with CDC as co-user of the CD 6600. New staff were hired to manage the freed-up capacity, and terminals were installed in the engineering institutes to make use of the new capacity. Users were separated from the machine and placed in front of IT terminals, sometimes right in the city at a very central location.

The growing demand for computing capacity in the two subcenters soon made further expansion of RUS necessary. The expansion was carried out in the "natural sciences center" in Stuttgart-Vaihingen, where the computing center had moved into rooms. By 1975, the organizational, "historically contingent division" of the computing center had already been superseded. Department B in the Stuttgart suburb of Vaihingen assumed sole control. In this expansion process, the CD 6600 was moved from the space in the Institute of Statics and Dynamics of Aerospace Structures to new space in the computing center. There it was hooked up to a brand-new CYBER 174, which was also manufactured by CDC. This move was "long overdue" in view of the "catastrophic undersupply of computing capacity at the University of Stuttgart," according to the rector's report. Finally, "interactive terminal operation" and thus "dialog traffic in remote data processing" could be introduced. The report referred almost euphorically to a "user-friendly variant to the conventional mode of batch processing," which "should lead to considerable time savings, especially in the development and testing of programs." Further expansion of the machine fleet by way of an additional CYBER 175 was included in the planning and requested from the DFG in 1977.[7]

Planning crisis and a flood of proposals

The center remained in flux and gained in operational importance. As machines became faster, the programs and users more numerous, expansion seemed only a matter of time. But these expectations would prove a disappointment. The late 1970s bore witness to another full-blown computing center crisis, and not only in Stuttgart.

Initially, expectations seemed to be dampened only by the construction schedule. "With the cancellation of the new computing center building, which was reaffirmed again this year, the provisional accommodation of the computing center in space not intended for this purpose is becoming a permanent disaster," reported the computing center urgently.[8] But the DFG reviewers tasked with assessing the carefully prepared proposal from Stuttgart to finance a CYBER 175 also signaled reservations about the expansion plans. The reviewers appear not even to have been convinced of the need to increase capacity.

In October 1977 a rather dismayed delegation from Stuttgart traveled to Aachen to see the DFG's chief reviewer to be briefed about "particularly thin" sections of their proposal. The reviewer's widely known enthusiasm for CDC did not alter the fact that he – like the other reviewers – doubted the evidence of need. Moreover, the "reasoning behind the selection of the CYBER 175" was unconvincing. Stuttgart would have to present alternatives. For the mainframe computer, a non-CDC variant along with a solution involving "distributed processing by means of medium-sized computers" should also be considered.[9]

This suggestion called into question the close link between the center's development and expansion of the existing computer system. Suddenly, sticking to expansion plans based on a hardware concept that relied on long-term continuity had become risky.[10] Stuttgart therefore supplied the DFG with graph after graph showing the development of demand and capacity use. But the choice of manufacturer was non-negotiable, because any change would have torpedoed the very essence of the expansion concept. The DFG stuck to its guns and in 1978 definitively rejected procurement of a CDC CYBER 175, which would have relieved pressure on the CD 6600.[11] In retrospect, not much remained of the "noble objectives" of "future IT planning," noted RUS director Karl-Gottfried Reinsch.[12]

No one knew better than Reinsch what that meant for the computing center. In the end, the required capacity was procured – not, however, according to the plans and goals of the computing center but rather "according to the user's own goals," as Reinsch tersely put it. Users covered their computing time requirements simply "by procuring small computers as part of their research projects," a practice that increased markedly from 1978.[13]

As a result, costs for maintenance shot up, and overcapacity at the institutes swelled. Worse, any attempt at central planning was time wasted.[14] Invoking the "concentration of computing capacity and expertise" in computing centers and praising the "organizational form" embodied by "large computing centers" as an "economical IT concept" was of little use. "Small is beautiful" was the order of the day and readily financed; any reference to the economies of scale of centralized provision of computing capacity was, consequently, obsolete.[15] Even inveterate insiders like the editor of *Rechenzentrum* magazine gloomily pronounced the imminent "death of the computing center." Given the "link between small computers and distributed processing," it would behoove the management of computing centers to "think hard about a new self-image."[16]

Attempted resuscitation

As director, Reinsch was unstinting in his efforts to wring funding from the federal government and the state to replace or relieve the old mainframe machine and to breathe new life into the computing center. Proposals were bulked up or, conversely, split into two parts. Subject matter was changed. And the information requested was supplemented with additional rationales, tables, drawings, reports and expert opinions.[17] Suppliers such as IBM and CDC were asked for comparable information about their machines. A benchmark report was ordered from the Max Planck Institute for Plasma Physics in Garching on the Cray-1 in operation there.[18]

Conceptual work was also proceeding apace, as evidenced by a 1979 document on supplying the University of Stuttgart with distributed computing capacity.[19] To be able to operate RUS as a central facility "in the age of distributed processing," it would have to abandon its previous role as a data factory and instead provide new services far beyond existing user groups at the institutes.[20] The question was what a university computing center of the future could do aside from efficiently running its users' programs through a powerful machine and delivering the results as quickly as possible. For Reinsch it was clear in any event that the variety of new and forward-looking areas of application would also determine the demand for additional computing capacity. "Deciding to work in new, future-proof, competitive areas in energy research, aerospace, medical imaging, chemistry, etc., will automatically generate demand for supercomputing capacity," wrote Reinsch in a commentary on aspects of information processing in Baden-Württemberg. The title – "What does it mean to operate a supercomputer, and to what end?" – was borrowed from Schiller. Replacing "universal history" with "supercomputer" had obviously been done for fun; replacing "study" by "operate" was shrewd. For operation implied procurement, and procurement was based on need.[21]

There were essentially three ways out of the crisis, three fields of action in which new functions and computing center tasks could be developed. First, the computing center of the future was uniquely positioned to assure stable connectivity between all kinds of computers at the university.[22] Second, it was also thinkable to seek major customers and investors outside the university. And third, an exclusive offer could be made for rarely needed peak performance at the limit of computability.

As the rector's annual reports for 1978–1980 show, the quest for a new role for RUS entailed a steep learning curve and consequently was not exactly what the university administration had in mind in ordering a change of strategy. The computing center continued to complain loudly to the administration about its capacity constraints and blamed the DFG for its current troubles. Reinsch likely hoped to be able to get the attention of the state ministry for science and arts indirectly. However, as he was himself a member of the computer user committee in the vice-rectorate for research, it can be assumed that, ultimately, he had to hear and shelve his own petition.[23] In the midst of palpable anger over the intransigence of the federal research funding bodies and the many rejected proposals, Stuttgart began to argue that RUS's future computing capacity would be of benefit to the state of Baden-Württemberg. Indeed, the installation of a "scientific supercomputer" was all the more urgent for the state now that the Stuttgart proposals had been rejected.

Was that a threat? – throwing our careful planning overboard will come at great cost – or did it indicate a change in strategy? The answer is probably both, as well as stubborn adherence to the model of the computing center as a data factory with a capacity problem. That is to say, the "urgent" procurement of a supercomputer was directly linked to the lack of mainframe capacity: "Tests have shown about 50 percent of our mainframe capacity could be advantageously transferred to this supercomputer."[24] In other words, Stuttgart was prepared to purchase a supercomputer to take the pressure off the mainframe machine.

That was easier said than done. It wasn't until March 1980 that a proposal for a scientific supercomputer for the state of Baden-Württemberg finally heralded the end of mainframe culture, at least rhetorically. This shift was evident in the prominently placed reference to the "consensus on funding policy" as stated in the DFG's 1979 recommendations on the procurement of data processing equipment for universities. In Stuttgart, the DFG recommendations were read and cited as a decision to continue to "address those classes of scientific problems that cannot be processed using conventional general-purpose computers or can only be processed uneconomically." Consequently, mainframe computers should no longer be upgraded or replaced. Rather, according to a quotation from the DFG's recommendations in the new proposal, boldfaced for emphasis, special computers should be installed in "regional computing centers."

Because Stuttgart had long been a "regional computing center,"[25] all that was actually needed was a "special computer" to fulfill the DFG's recommendations. If required, specific offers could be produced (for example, from Cray, CDC, and Burroughs). "In view of the country's still open procedural approach," however, no detailed analysis of bids would yet be made. The awkwardly phrased "open procedural approach" spoke volumes. No one was eager to give the impression that, once again, everything had already been predetermined and finalized in Stuttgart. Accordingly, details of hardware and software, and of administrative guidelines, would be left to "higher-level committees."[26]

Unfortunately, the DFG's "consensus on funding policy" also affected the regional program, with immediate implications for Stuttgart. The program ended in 1981. Thereafter, the DFG almost completely ceased to fund computer installations for research projects, since the purpose of computing was now considered to be "ensuring capacity." In the future, federal funds for large computer installations would only be available through the University Construction Act (HBFG).[27] Consequently, the University of Stuttgart sought to acquire a computing center building with a cutting-edge vector computer through this source of funding.

The fact that the University of Karlsruhe still submitted its proposal to procure a vector computer under the federal data processing program posed no problem for Stuttgart.[28] As a neighboring university, Karlsruhe was naturally a competitor. But the HBFG approach Stuttgart was following in the meantime had its own advantages and seemed to alleviate the pressure. The simultaneous southern German grab for one nearly empty and one half-full federal funding pot would not in principle have come as a surprise to anyone in Bonn. In the late phase of West Germany's cooperative science policy, it was calmly and quietly noted that a state-specific need for coordination was becoming manifest via the federal capital. Both proposals were duly assessed by the DFG and rejected. The message was clear: prior to making any purchases whatsoever, the state was going to have to discuss actual need.

The Baden-Württemberg rector's conference thus was compelled to set up a working group to keep the conflict over scarce resources at a controllable level and to work out a defensible consensus to be resubmitted "to the federal government."[29]

Building consensus consumes time and effort. The chairmen of all the senate committees tasked with overseeing computing centers in Baden-Württemberg's universities, and all the directors who ran those computing centers, had to sit down together and make "recommendations for the selection of a vector computer, operating modalities, and, of course, the location."[30]

The recommendations developed gave equal coverage to all three ways out of the computing center crisis – expanding the customer base, dealing with communication technology problems, and focusing on computer operations. Karlsruhe and Stuttgart would each have to reduce their proposals for a vector computer to half of the procurement cost. Both universities compensated for the resulting financing gap by expanding their customer base for vector computers, including customers outside the university. Industry contributed to the capital expenditure costs, the universities ensured the operation, and the computing centers established the necessary connections, including to the proposed specialist computers.[31]

Thus did RUS emancipate itself from mainframe computing after many years. Despite the curtailing of federal government

zurück

subsidies, Stuttgart managed to squeeze into the exclusive club of supercomputer owners through the acquisition of a vector computer, which, being a discontinued model, was purchased at a substantial discount from Cray.[32]

Shaping policy and organizational structure

In fall 1983, the Cray-1 was installed in the space of the university's old central kitchen in Vaihingen.[33] For theoretical chemistry, whose slogan was "Don't cook, compute!", the address was certainly appropriate. But the machine also stoked the confidence of science policymakers in Baden-Württemberg. In June 1985, during a trip to the United States he undertook together with the rector of the University of Stuttgart, Minister-President Lothar Späth casually ordered in Minneapolis a Cray-2 – without a university discount, a loan from the state, or prior proof of performance. The financing could be sorted out later. In this way, Stuttgart secured the Europe's first Cray-2, the best of the best, a computer that would, naturally, be able to calculate much faster and that was possessed of an especially large memory.

Even before this magnificent machine was delivered and put into operation, it changed the realm of expectations at RUS. The computing center had begun to free itself from the federal government's coordination framework and to adopt a new modus operandi. Instead of relying on the skillful implementation and combining of existing plans, or considering the coordination efforts of the federal and state governments as a precondition for its own actions, Stuttgart began working on the agenda. Späth wrote of a "pivot to the future," stressed the malleability of circumstances, and pursued an innovative mix of science and economic policy.[34]

In view of the new machine's performance class, and the new competitive trend in policy, the university wished to consider a few ticklish questions pending delivery of the Cray-2.[35] In the process, a new openness became apparent. Here, too, circumstances were discussed in terms of their malleability. An expert seminar on the applications, financing, and organization of supercomputing clarified the "sometimes widely divergent points of view" of the experts and confirmed "the necessity of addressing the subject." Administrators and financial experts had to learn to appreciate what supercomputers were for and to "come up with new ways of financing and organizing procurement and 'distribution' by dint of their own conviction and imagination," wrote Jürgen Blum, chancellor of the University of Stuttgart, in the foreword to a report on the seminar.[36] The seminar program traced an elegant arc. The first three speakers were from RUS, Forschungszentrum Jülich, and Adam Opel AG in Rüsselsheim. Additional contributions were heard from the Konrad Zuse Center in Berlin, the Institute for Business Administration in Erlangen, and finally (and in detail) once again from Blum. The organizers thus succeeded in covering a broad swath of current issues. Those who experienced the seminar also had a good grasp of what running a high-performance computing center entailed.

Roland Rühle, RUS's new scientific director, kicked off the event, assuring the participants that there were indeed many conceivable applications for a vector computer at the university. These applications were also "extremely diverse," he said, because they comprised anything related to problems of high dimensionality, complexity, ill-conditioning, and nonlinearity – from nonlinear structure-mechanics analysis, fluid mechanics, reaction kinetics, nuclear technology, and plasma physics to molecular structures and binding energies. However, efficient numerical methods would first need to be developed; and that meant integrating software from different computers with different operating systems and compilers. Also on Rühle's wish list were consistent user interfaces, a database system, suitable graphics equipment, a powerful computer network, and well-trained users. That was because operating a supercomputer required "synthesis of knowledge from the fields of applications, computer science, numerics, and computers."[37]

Following Rühle's overview, Rolf Theenhaus of the Jülich Nuclear Research Facility (KFA) discussed supercomputing at large-scale research centers. In Jülich, he said, it was no longer necessary to explain the reasons for operating a vector computer. In this case, the experimental facility and the computing infrastructure were already well coordinated. Consequently, he had the luxury of describing his institute's excellent situation. He began by emphasizing the importance of the thirteen large-scale research centers in Germany in general. He then described the Jülich KFA in particular. Finally, he provided a detailed explanation of how vector computers had replaced the sequential architecture of John von Neumann's universal computers. No sooner had Theenhaus embarked on vector computer applications in Jülich, however, than the dreaded occurred: in discussing his first example, the Czochralski method of growing crystals, he lost his audience. His second example, "surface phenomena," may have roused the seminar attendees, but only up until he reached the following complex phrase: "The density-functional formalism (Hohenberg, Kohn, Sham) provides a suitable method for calculating the ground-state energy of such a system, thereby simplifying the original many-particle Schrödinger equation to a one-particle Schrödinger equation, for which a self-consistent potential must be determined by solving the Poisson equation."[38]

Werner Martin, director of engineering and manufacturing systems planning at Electronic Data Systems Europe, spoke next.

Contrary to expectations, he said, a supercomputer could also play an important role in the real world, for example, in designing cars at Adam Opel AG. Creating everyday industrial products relied on computer-aided logistics and processes. At the same time, developing a car for the mass market required highly computationally intensive flow calculations and simulations, for example, in crash tests. Opel had to solve both conventional computational problems and tasks requiring supercomputing, and to that end employed two different types of machines that interacted with each other. For RUS, the situation was a familiar one, as the Cray-1 had been connected to front-end computers and the old CD 6600 – dared one say it? – was still running. Opel relied on a mix of equipment because the tasks to be solved in the automotive industry were especially heterogeneous. In this way, Martin succeeded in addressing not only the audience from RUS and the university administration but also the next two speakers, who dealt with questions of organizing and financing the Konrad Zuse Center for Information Technology Berlin (Peter E. Schuhe) and with distributing the investment and operating costs of supercomputers in the context of cooperation between industry, large research institutions, and universities (Wolfgang Männel).[39]

The contribution from business would have been particularly gratifying to Blum as chancellor. Männel explained that the acquisition of supercomputers for the university would ensure extremely efficient, large-scale processing. In principle, this should result in a very favorable “cost–performance ratio,” because “[the computers’] capacity is such that it leads to significant cost degression.” In view of the high fixed costs, however, this advantage would only be realized if the university operating a supercomputer cooperated with other universities or with large research centers, thereby creating conditions for “the user group as a whole to properly exploit the considerable computing capacity.” This was patently obvious, but it

zurück

forced a rethink on Stuttgart's part: first, it would be necessary to get used to the idea that in future, RUS would have to deal with the problem of overcapacity rather than of capacity bottlenecks; and second, “cooperation among users,” according to Männel, “inevitably implied the question of how costs are to be allocated.” Moreover, distributing costs would in turn make cost transparency necessary. “The exact definition, differentiation, tracking, and documentation of all costs incurred by the provision and use of supercomputers is by no means a trivial problem,” said Männel, who proceeded to detail evidence for his claim. Probably to the surprise of those present, what ensued was no mere rote exercise but rather a knowledgeable discussion of the advantages and disadvantages of different forms of billing. Should the transfer price be calculated using full-cost or marginal cost accounting? Could scarcity pricing also be used as a guide, or was there something akin to a market price? The answers to these questions had economic consequences, but ultimately could only be decided by clarifying the desired research, capacity, and cooperation policies.

Männel thus passed the ball back to the university administration, in the form of Blum, who gave the last talk. As a lawyer, and soon-to-be doctor of administrative sciences, Blum was interested in fundamental questions of university governance. The computing center was a particularly exciting example because the sheer size of the capital expenditure involved was a challenge for conventional line-item budgeting. How could an investment of 80 million marks, involving state and private parties, be fitted into a university budget that was based on annual allocation of funds and monitoring of current expenditures? Moreover, how was this investment to be legitimized when it was made ahead of demand, that is, when computing capacity had to be pitched to and placed with a very heterogeneous user community? Blum expected a self-regulating incentive and control system, and he saw the need for user targeting – which might jeopardize the decision-making power of the center. Nevertheless, he said, now was the time to dare venturing “to the edge of legality” in choosing an organizational model. Policymakers, Blum stressed, were urging competition. Efficient organizational models would certainly be rewarded, whereas weak ones would be scuttled. Would conventional business models be suited to reconciling “autonomous use of resources” and protracted capital investment processes?

For now, these were just hypothetical questions. But they indicated a shift in the strategic focus of RUS and the university administration. Going forward, they did not wish to have to align everything according to the rules of the DFG, the state, and the federal government, nor did they wish to continue submitting proposals in the face of conflicting priorities. If new rules were needed, they wanted to help shape the local implementation. Finally, the importance of university computing centers had changed in the meanwhile, even in terms of their centrality. “The installation of mainframe computers and supercomputers in a central university computing center,” said Blum, “does not imply the centralization of data processing capacities. In addition to centrally operated mainframe and supercomputers, there are an increasing number of decentralized powerful computers in research institutes that are networked with each other and with the computers in the computing center and which, thus, also have access to national and international computing networks.”[40]

A diversity of machines

Making proper use of supercomputers requires drastically freeing up their resources.[41] The purpose of a Cray-2 is “solely to utilize computing power,” stated RUS in a brochure about the new supercomputer. “All other services, such as file servers with archiving, dialog servers, graphics servers, print servers, and network servers must be provided by front-end computers.”[42] That meant that, in Stuttgart, the Cray-2 would have to interact with a considerable number of other machines in the computing center itself as well as in the research institutes associated with the university. The fact that computing centers operated a diverse collection of machines was not unusual, and not particularly worth mentioning. In describing themselves, computing centers typically highlighted brand-new machines or very powerful ones. Equipment that had already been in operation for a long time, or that merely prepared input and evaluated output, was simply expected to perform the designated service without any special distinction. What was special about RUS, which installed the Cray-2 in autumn 1986, was not the diversity of the existing machines. What was new was that this mix became part of the public identity of RUS.

Despite its neat rendering, a graphic depiction of RUS’s machine fleet from 1986 was still fairly confusing (Fig. 2). The many boxes labeled with computer names and types could barely be deciphered without accompanying text. Not because the designation “IMB 3083 J16” was meaningful at best only to specialists. Rather, the difficulties of representation and interpretation had to do with the fact that individual machines could not be assigned a single function: “In Stuttgart, a CYBER 835, an IBM 3083, and a VAX 11/780 are currently being used as front-end computers.” At the time, a front-end computer performed several functions. Moreover, although one could rightfully say that the “dialogue and file server functions are being largely shifted to the IBM” and “the network server functions are being taken over by the VAX,” there were three VAX units. In addition, there were two MicroVAX computers and no fewer than five other VAX’s located at the research institutes, whereas the

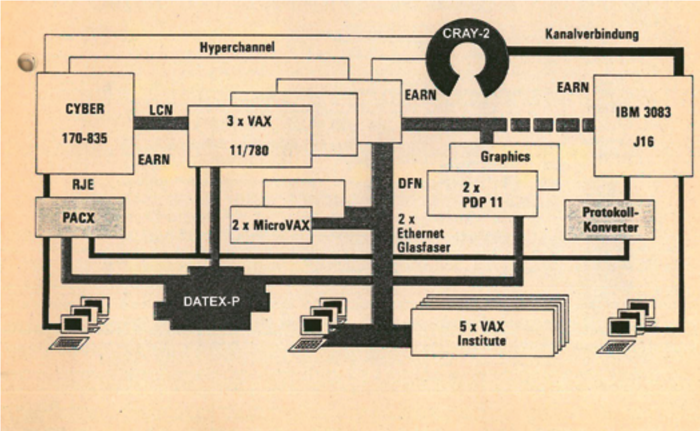

Fig. 2: Type and disposition of machines in the computing center, 1986.

much older PDP 11's were apparently closely coupled to the very modern graphics computer.[43]

The schematic simultaneously served as map, agenda, inventory, and concept. It could be interpreted as an overview, a functional diagram, a historical reference, and a horizon of expectations, with descriptive textual accompaniment.[44] Assigning several functions to a front-end computer had already proven worthwhile with the Cray-1. Consequently, the brochure noted with respect to the functional diversity of the machines that with the "exception of some components," the same visual convention was "to be retained for the CRAY-2 as well." Computers that were soon to be replaced were missing from the schematic, whereas others, which had not yet been installed, were included. For example, the horseshoe at the top of the schematic was already labeled Cray-2, although the graphically absent Cray-1 would continue to run for many months. The Cray-2 had yet to deliver on the promise of its purchase.

The graphical tidiness of the schematic as well as its functional and organizational compression – also reflected in the staggered arrangement of the workstations and terminals – created space to associate machine diversity with an impressive network diversity. “Embedding a supercomputer in a computer network is of paramount importance.” And, indeed, the brochure listed virtually all of the network protocols and cables known around the mid-1980s. These were the prerequisite for “simulation and interactive applications at the workplace.” The entire arrangement worked very well at the Vaihingen campus thanks to the “fiber-optic ETHERNET” connection “with the VAX from CRAY.” The VAX computers and DECnet were linked to the terminals and graphic workstations and connected them with all the other workstations and “the CRAY.” Exactly which Cray was not specified, and the text segued quickly on to general plans for the future: adding simpler workstations and PCs to the network.

Nonetheless, the brochure called access to the supercomputer “from outside the campus” problematic. There was a wide range of options – switched lines, dedicated lines, the IBM research network EARN and the German Research Network via Datex-P. But these networks all delivered substandard performance and represented expensive emergency solutions. As a rule, connection to the supercomputer from outside did not work as well as on the RUS-served campus in Stuttgart-Vaihingen. RUS was quick to delegate this task to an ISDN pilot project run by the federal post office (Bundespost) in Stuttgart, which began in 1986 and which the university also wanted to employ to test “its computer–computer connections, in particular between workstation computers and the CRAY.” The state government of Baden-Württemberg, in cooperation with the Bundespost, intended to “make fiber-optic connections available primarily for industry and research," as was already possible in the Vaihingen university area.

The project enabled RUS to demonstrate its competence in dealing with a variety of apparatus and procedures. The computer landscape was depicted not as a confusing cabinet of curiosities and a product of historical accident but as a powerful tool for dealing with complex capacity issues. This tool assured the supercomputer's performance by means of an arsenal of front-end computers, data servers, and special computers for displaying results and provided innumerable interfaces to connect the center's own and external computers "with the Cray."

The diversity of machines and networks in turn enabled a heterogeneous user community to deal with a wide range of very different applications. RUS was no longer just a locus of centralized resources with homogeneous operating rules. It was a machine and application cluster that counted on the heterogeneity of its resources. It built bridges between these resources and adapted to the thematic and disciplinary variety of its users' tasks (Fig. 3). What made RUS a center was that it was able to achieve this integrative capacity through the diversity of its machines and networks.[45]

Shielding users from complexity

The organizational shape shifting, the array of operating options, and the highly networked technical heterogeneity of a "computing center with a supercomputer" could be tamed by "imaginatively" venturing "to the edge of legality" (as Blum put it) or by embracing resource diversity (according to Reinsch and Rühle). From the vantage of the operators, many of RUS's problems around 1985 were issues that could in principle be solved. In the

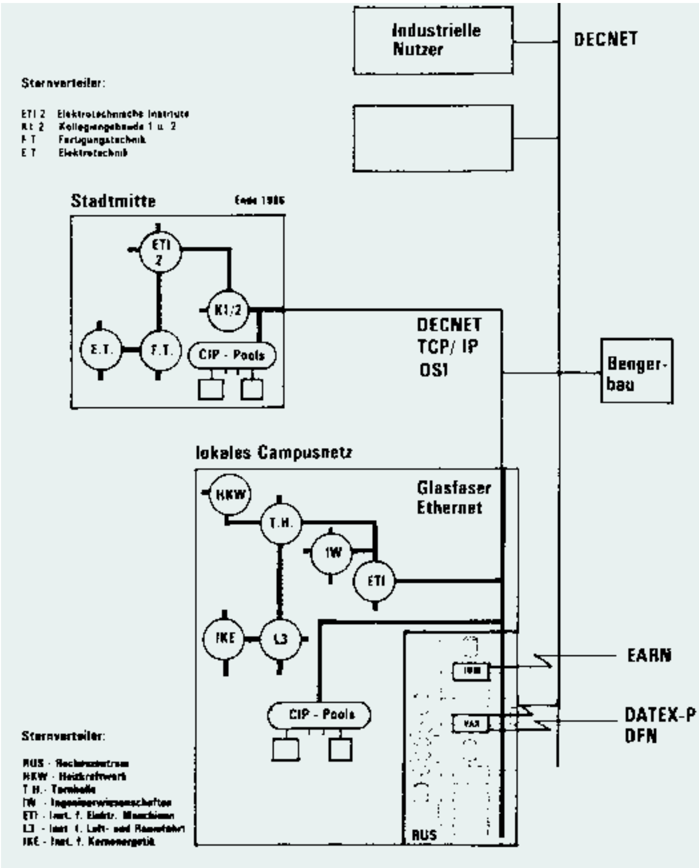

Fig. 3: The integrative capacity of RUS, 1987: many protocols, distinct locations, powerful computers.

computing center, for example, the operators knew the meaning of cooperative distribution of computing capacity and sustainable payment terms.

For users, however, getting a grip on the complexity of the computer and connection clusters was more difficult. The computing center had to try to make the differences between programs, data, networks, computers, tasks, evaluations, and disciplines disappear so as to shield users from the heterogeneity of the actual computing center, which only got in the way. For a user, the technological divide between an IBM computer and a vector computer from Cray was better breached if it could be ignored or not seen at all.

If supercomputing was to be carried out from workstations at the institutes as well, additional navigation aids would be imperative. In a project report from late 1987, Isabel Loebich, Lothar Ehnis, Ulrich Lang, and Roland Rühle stated: “Distributed computing in heterogeneous computer networks comprising workstations, supercomputers, graphics devices, and laser printers often requires the user to have knowledge of several operating systems and transmission routes.” Distributed applications in heterogeneous computer networks provided a way for RUS both to mediate between protocols, languages, programs, and systems and “to shield users from this heterogeneity.”[46]

The problem had long been a familiar one to RUS’s scientific director. As early as the 1970s, Rühle had been working on a tool that would enable “knowledge in the form of algorithms and data to be [stored and] retrieved at any time” and that could “reduce the large number of possible interfaces to common invariant elements through abstraction.” The result was RSYST (*Reaktorsystem*), which already in the heyday of mainframe culture and entirely in the service of reactor safety provided a bridge that, aside from its “primary goal of integrating programs and data,” enabled "integration of different disciplinary methods."[47] The machines, programs, and networks that populated the vast terrain between supercomputers and workstation computers at RUS in the mid-1980s made RSYST still an attractive tool for dealing with heterogeneous computing. RUS lost no time in adapting its successful navigation tool to newer developments and to integrate it, for example, into the protocol requirements for the German Research Network (DFN). "DFN-RSYST" provided "the user with a uniform view of the network." It enabled "the arbitrary distribution of pre- and post-processing, computationally intensive tasks, and – especially graphical – network analysis" by providing "mechanisms for outsourcing calculations, transferring data, and accessing data via the network."[48]

Communicating to the public

In the mid-1980s, RUS underwent a profound change. New computers were purchased, old computers were assigned new tasks, and the research institutes were integrated into the center's operations as part of a distributed computing scheme. In addition, mainframe-assisted data processing was transferred to a supply-oriented competence center, which reigned over a very heterogeneous ensemble of applications, interfaces, programs, and machines.

This change demanded a communications tour de force on the part of the center's operational and scientific management. RUS's machine fleet was hardly self-explanatory. Moreover, owing in part to the intervention of the minister-president, it had become a surprisingly public and political affair.

This situation was detailed in an extensive article in the March 1987 issue of *Computer Zeitung* magazine. Since the first

Fig. 4: Skeptical audience, 1986: the minister-president defends his pivot to the future.

CeBIT computer expo in 1970, the magazine had provided a broad IT readership with well-researched articles and analysis on a weekly basis. Now it was Stuttgart's turn. The focus of journalist Ludger Schmitz's article was immediately evident from the title: "The University of Stuttgart computing center acquired a Cray-2. Then the trouble started."[49]

The article delivered a wide-ranging critique. It began by referring to the justification for the state's procurement of the Cray-2 as "somewhat misleading," despite the unexpected good fortune of a falling dollar exchange rate. In purchasing the computer, Späth had ignored the "official channels" prescribed by the federal government, which prompted the DFG to ask the Science Council (*Wissenschaftsrat*) to investigate. The resulting report was "scathing in its criticism." No test runs had been made, and no university discount had been negotiated. "Instead, Späth used the University of Stuttgart's privileged status to embarrass other German universities and research institutes." According to *Computer Zeitung*, citing the Science Council, the contract contained a clause preventing Cray from delivering the same type of computer to another German university for one year.[50] The result was to "hinder research," "artificially create demand for computers" in Stuttgart, and distort "scientific competition."[51]

The magazine's litany of accusations was not only directed against Späth (Fig. 4). There was also the matter of the substantial difficulty involved in getting the Cray-2 up and running.[52] Moreover, it was no secret in the community that the inspection report for the Cray-2 "was a sham" and that the sale of computing capacity to industry and to research facilities was slow.[53] Although the fiber-optic cables for external users had been laid, the connections were still missing because pricing had not been agreed with the Bundespost, making interactive operation on the Cray-2 unaffordable. Simply lowering the price of computing time for external users in response to lack of demand was out of the question, since it would create the appearance of distorting the market through industry subsidies. *Computer Zeitung* was particularly critical of the commitment made with the purchase of the Cray-2 to prevent any use by countries of the Eastern Bloc. The US government's export restrictions were found among the exclusion criteria on the application forms for granting a user number at RUS. *Computer Zeitung* obligingly printed the form.[54]

To the computing center, the university, and the state government, it felt like the sharks were circling. Späth may have been clever enough to circumvent the federal government's funding rules, but it also stood to reason that the DFG could retaliate through the Science Council. *Computer Zeitung* had gotten wind of precisely that and accordingly painted readers a picture of the dog-eat-dog world of supercomputing. Policymakers were urging competition, Blum had stated at the expert seminar. To be competitive, assets had to be sufficiently well protected through appropriate measures. Conspicuous science and research policy had to be supported by effective public relations. RUS urgently needed to find ways to explain the Cray-2, because the Cray-3 was sure to follow.

Consequently, in the wake of the Cray-2 procurement, Karl-Gottfried Reinsch and Roland Rühle found themselves having to pay close attention to the public, who were closely following events at RUS. "Supercomputers are currently a hot topic. Daily press and trade journals do not always do a good job of reporting on these developments," complained Reinsch in 1987 in a detailed article that called the Cray-2 a "challenge for science and industry." That this opening salvo was unlikely to win over the daily press and trade journals was evident.[55]

Reinsch attempted to approach the subject systematically. What were the requirements in terms of memory size and computing capacity for different areas of application? Where did Stuttgart stand in the "league" following the acquisition of the Cray-2? Which users from which scientific domains required which amounts of computing capacity and which programs? How did one go about getting one's project "into" the supercomputer? And where did this machine fit within the University of Stuttgart's computing infrastructure?

To answer these questions, Reinsch resorted to a whole range of graphic forms of representation – plunging himself and his audience into a communicative abyss. On closer inspection, the three-dimensional schematic showing the relationship between memory, speed, and problem size was actually comprehensible. But the multicolumn table comparing computers and manufacturers was largely unreadable; and the hierarchical-pyramidal ranking of supercomputers could only be understood with the aid of the accompanying explanation. According to Reinsch, the functional "integration of the supercomputer into the concept of computer use" at the University of Stuttgart could be illustrated by the "4-level model." What a reader was supposed to deduce from the detailed hardware descriptions – including weight, floor loading, coolant temperature, and closed-circuit coolant flow – was unclear, despite the intended transparency.

Reinsch proceeded to give examples of the types of problems the supercomputer could process, among them "the structure of the thyroid hormone thyroxine," the "density distribution of interstellar gas in colliding galaxies," and (naturally!) a "numerical simulation of the three-dimensional development of vortex structures in the laminar-turbulent transition in boundary layers" – all supplied as figures. Additional projects were described in the main text. The conclusion reiterated that supercomputing was very expensive, but that with increasing use it would likely become "possible for everyone."

Roland Rühle's approach to public relations was totally unlike Reinsch's. Rühle was convinced that the proposals researchers sent to funding agencies were ill suited to publication. Journalists would be very unlikely to pick up the material there and make it available to a broader public in a reader-friendly way. Accordingly, when Rühle spoke to the *Industrie-Anzeiger* trade paper, his intent was not to announce but to communicate. The information should be tangible and easy to explain. The message should be neither annoyance (*Computer Zeitung*) nor a "challenge to science and industry," (*Praxis der Informationsverarbeitung und Kommunikation*), but rather a "tool for engineers: the world's fastest computer calls the University of Stuttgart home" (*Industrie-Anzeiger*). This article, clearly supplied by Rühle, ranged from the Cray-2 and the new ways in which Baden-Württemberg was promoting high technology to the network of universities, research institutes, and companies such as BASF, Hoechst, Bosch, Porsche, Dornier, TÜV, and Daimler-Benz. The prospect of the imminent use of the supercomputer by small and medium-sized enterprises (SMEs) was also hinted at.

Jürgen H. Koch, a journalist from Munich who had been briefed by Rühle, proved to be an outstanding communicator. That much was clear from the references to "working together," "downtown," a "pilot project" with the Bundespost, Karlsruhe and Freiburg, and, of course, "fiber-optic cables." Koch fearlessly reported and compared costs for computing time, cited cooperative arrangements with private companies, and touted solutions to complex problems. None of it was trivial. Nevertheless, an elegant, politically inclusive way of communicating about the Cray-2 made it easier to grab the attention of readers. Because what the Cray-2 was, what would result from it, and what might follow it would all have to be discussed again in the foreseeable future.

The performance gambit (1988–1996)

Supercomputers are difficult to understand – not only because they calculate at what seem to be near-miraculous speeds. Even more difficult is figuring out how to optimally configure the machine and everything connected to it. A complex ensemble of computers, institutes, and ministry officials must be able to work together; a fragmented committee landscape has to be made navigable; and the different interests of computer manufacturers and scientific and industrial customers of computing capacity must be reconciled. This is easier to do when evaluation is based on consensus. Technical support, fast access, user autonomy, attractive pricing, efficient programs, simple pre- and post-processing, and package solutions are all considered. That a supercomputer should be able to provide "unrivalled performance" hardly needs saying.[56] What is surprising is how elusive this criterion can be when it comes to daily operations.

Procuring the Cray-2 was the deus ex machina that resolved the hopeless tangle that the Stuttgart computing drama had become. No years of haggling. No cooperation problems. Nobody to get in the way. Everything happened much too fast for that – with respect to the purchase as well as the computing capability. Only afterward did the tedious work of configuration begin. In the process, many contradictions emerged. First, Stuttgart enjoyed an exclusive facility, but at the same time had to expand its user base. Second, Stuttgart could not avoid dealing immediately and very specifically with the next replacement procurement. Only in this way would it be possible to keep up with computer performance. That meant having to plead for a computer called the Cray-3 that

zurück

did not yet exist and, ultimately, never would exist. Third, evaluating computers began to be easier thanks to a consensus-based benchmark that was gaining acceptance in the field of supercomputing. The benchmark served as a powerful advocate when it came to procurement, but only if machines could be tested and manufacturers could provide useful data. Which could hardly be expected in the case of future facilities. Fourth, like other computing centers, Stuttgart struggled with architectural imponderables and inconsistencies. Which was the better choice? Tried-and-true vector computers, parallel computers, or a network of workstations?

Despite the highly autonomous solution in the case of the Cray-2, the scientific and political framework also imposed contradictory requirements for the current and future configuration. What would be the cost if a national supercomputing center had to guarantee access for the entire country yet was predominantly dependent on regional funding? In Stuttgart, some of these contradictions were resolved with an original but complicated organizational structure that separated capital investment and allocation of computing time but that required combining business and science. This may explain the gradual substitution of the notion of supercomputing by the more operative term "high-performance computing."

Simulation for all

Enabling a global player like Porsche to profit from the efficiency of the Cray-2 required little in the way of justification. Simulations on fast computers were nothing new for Porsche's engineers. They had been simulating the crash behavior of their sports cars for a long time and could dispense with actual dummies colliding

Fig. 5: Crash and simulation in the 1990s: a Porsche 911 Carrera after a frontal offset collision.

with actual airbags and steering wheels.[57] The Cray-2's processor networks now even allowed the engineers at Porsche to calculate their car bodies with ever finer resolution. It was already possible to simulate the first 80 milliseconds after impact. The engineers followed the events on a computer screen, optionally in slow motion, as often as they wanted (Fig. 5).[58]

Harder to justify was the intention to turn even SMEs into Stuttgart supercomputer users. The sheer availability of computing capacity at the top end of the market did little to motivate companies to get seriously involved with supercomputing or simulation.[59] The staffing levels in SMEs were far too thin, and the opportunity costs far too high. Any alignment of the supercomputer in Stuttgart and the potential computing requirements of SMEs would have to be planned and looked into by others. In a politically concerted action, the state of Baden-Württemberg, the city and University of Freiburg, and the Chamber of Industry and Commerce Southern Upper Rhine together founded the Freiburg regional enterprise High Tech Computerdienste Oberrhein GmbH (HTCO).[60] The goal of the joint venture was to explain to SMEs the link between computing excellence and regional economic development. HTCO expected no miracles and figured "that it [would] be pretty tough going at first."[61] The forty-five potential customers between Basel and Offenburg would have to learn the "basic principles of supercomputing," which was no mean feat. The only easy part was the message communicated by the distributer via the *Badische Zeitung* to companies in southern Baden-Württemberg: supercomputing was another word for simulation. It was only "thanks to these extremely fast computers" that "the idea of computer simulation" (which was hardly new) had become attractive to SMEs. The editor also had the numbers to back his case: "What it took a conventional device 18 *hours* to do, the Cray-II manages in two *minutes*."[62]

Two minutes instead of eighteen hours. Despite all the rapid-fire demonstrations, the explanation of supercomputing must have perplexed at least one or two SME managers. In contrast to other computer applications, they were told, high-performance computers were not tools for rationalization. Computers like the Cray-2 were not simply calculating machines to throw numerical data at. "We don't want customers ringing the doorbell and saying, 'Now do the math,'"[63] HTCO manager Franz Heidinger told the *Badische Zeitung*. Rather than just provide data processing for South Baden firms, the supercomputer experts from Stuttgart wanted to advise people and to find out "whether and how specific development and optimization problems can be solved using the

zurück

Cray II."[64] Exactly what this advice and problem solving were intended to achieve by means of simulation was still not very clear, despite Heidinger's references to an awning manufacturer "who wants to construct a support arm for his awnings that is three times lighter but twice as strong."[65] What was not under dispute was HTCO's basic premise: "The principle of high-performance computing can be applied to all optimization and simulation tasks."

HTCO's marketing strategy was unable to resolve the paradox between the exclusivity of the high-performance computing offer and the widest possible group of users. Moreover, the imagination of the potential user community was focused on the Stuttgart supercomputer, as if the mere presence of the Cray-2 in Stuttgart was enough to convince SMEs to delegate their research and development work to computer simulations, and thus enter the high-tech age. In any case, articles in the local press concerning the activities of HTCO pictured only the semicircular, human-sized Cray-2.

In contrast to regional business development, RUS did its utmost to reinforce the exclusivity of Stuttgart's computing services at the limit of computability. However, the means available to the computing center for symbolic capital were limited. Advertising the arrival of the Cray-2 would hardly suffice, as Reinsch and Rühle had learned from their *Computer Zeitung* experience. And neither computer simulation nor development of aesthetically appealing, interactive user interfaces was a special feature of the Stuttgart services that could only be taken advantage of there and not, for example, in Berlin, Garching, or Jülich. Consequently, RUS began to combine simulation and visualization, a new area of endeavor that was otherwise offered only in the United States (Fig. 6).

Simulated images constituted a production of Stuttgart computers performing at the limit of computability.[66] They also

zurück

influenced the semantics. The Cray-2 would open “a completely new dimension of science and technology simulation.” In place of “endless, barely comprehensible columns of numbers,” there were now “visualizations of technical procedures and processes as well as representations of scientific models.”[67] The Cray-2 was no longer just a supercomputer but a producer of artificial evidence. The simulated images generated by the Cray-2 promised insights into the physical world even though they were created in the computer. These images made tangible an artificial world thanks purely to supercomputers and visible only with the aid of monitors and printers.

The cost of visualizing computing output

RUS began to use the new tangibility of simulation for purposes of communication.[68] Simulated images were liberally incorporated into lectures, applications, and publications. One did not even need to know which double-star system problem was of interest to the theoretical astrophysics department of the University of Tübingen or why the University of Stuttgart’s Institute of Mechanics was preoccupied with airflow over a car sunroof to be fascinated by these new types of images. The aesthetics of the Cray-2’s housing may have been elegant, and the tabulated benchmarks of the supercomputer impressive. But compared with the new simulated images, both were much too mundane to enable people to “see” the supercomputer’s performance.

Visualizing mathematically simulated conditions required a great deal of work. For one thing, users had to be integrated into the simulation processes and thus also into the infrastructure of the computing center. But, in fact, the work began even before the simulation was completed. Even the definition of the problem

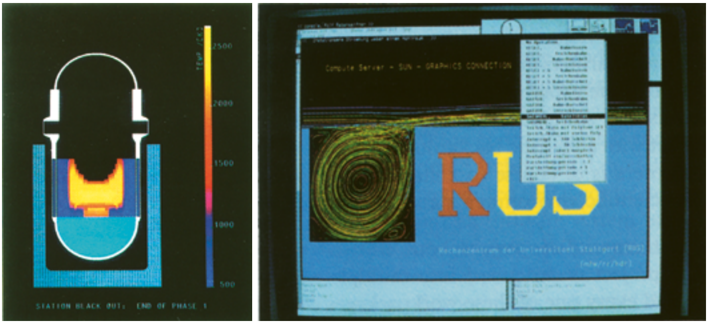

Fig. 6: Visual simulations, 1993: powerful images for human users.

had to be divided up by users in such a way that it could be handled in modules by different computers and programs and appropriate points for user intervention could be indicated. Only then could the modules be integrated as parts of programs and data sets into the RSYST in-house operating system. Since not every "decomposition of the task into sub-problems" resulted in "numerically satisfactory simulations," modular decomposition and numerical methods had to be coordinated.[69] For users, the task at this step was to ensure the "combinability" of the individual modules and to map the "simulation tasks as flexibly as possible to the integrated program system."[70]

However, the attractive power of the images and the urgently needed interaction of users with the ongoing simulation processes raised a problem that cultural critics lamented and that posed an interesting challenge for supercomputing centers: the problem of "data overload."[71] The more complex the definition of the problem, the greater the amount of data. Technical simulations from the fields of engineering and natural sciences running

zurück

on the Cray-2 produced “output data in the order of 100 to 1,000 million numbers.”[72] Neither sophisticated algorithms nor efficient operating systems could really deal with such large volumes of data. Hardware was also needed to simulate complex applications. In the computing center, the problem of burgeoning data volumes was translated into a task for infrastructure. Finally, the problem of data transport “from the supercomputer to the corresponding file server and, later, to the display terminal for analysis” had to be solved.[73]

The first bottleneck in this design was the file server. The server needed to provide high computing power and high bandwidth, but also have sufficient storage capacity to enable data to be buffered before sending it to the end devices. For this problem, the computing center reverted to tried-and-tested solutions. The long-standing connection to the manufacturer (Cray) was activated. And, probably to the surprise of the German user community, another supercomputer – the Cray Y-MP – was installed as a “high-performance file server.”[74] The public showcasing of RUS’s infrastructure policy was also new. The Cray Y-MP was commissioned “in May 1991 in the presence of minister [Klaus] von Trotha and the rector.” Finally, according to the rector’s statement of accounts, Stuttgart had succeeded “for the first time in Europe, and as the second institution worldwide, in procuring and successfully operating a file server in this performance class.”[75]

The second bottleneck in dealing with the flood of data in the area of simulation was the user seated before the screen. Here, the center’s computer scientists had to proceed more delicately. They employed contemporary scientific theoretical language and modeled users as a neural information processing system. “Data analysis,” in the words of the RUS simulation department, “is still determined by human perceptual capabilities.”[76] Humans are, and remain, “the slowest link in the chain of processing.”[77]

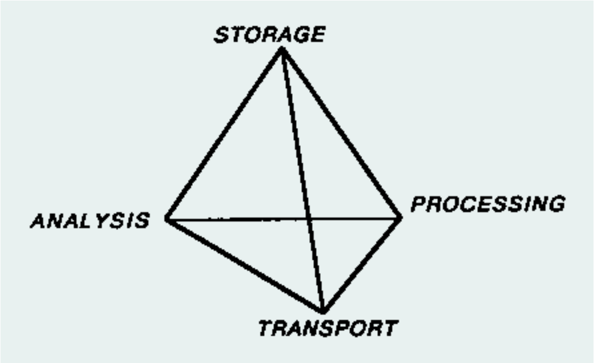

Fig. 7: Simplifying high performance, 1993: the Stuttgart simulation tetrahedron.

"[E]xploiting new findings on the processing of visual stimuli in the human brain" would help users learn how to acquire information much more quickly."[78] As far as incorporating human users into the local simulation environment went, the focus was "primarily on unconscious processing capabilities," which "do not require intellectual input."[79] It was assumed that the eye and brain together represented the "most efficient channel for transmitting information in humans."[80] Because humans are particularly good at unconsciously recognizing patterns, textures, contours, and movement,[81] users should be integrated into the simulation process "through the use of graphic aids." The deficient human, who could only understand the volume of data as an image, was placed in a trade-off relationship with the no less deficient computer system, which had to be given detailed instructions, data, and problems. Once the users were connected to the local systems, the computer scientists at the computing

zurück

center no longer needed to look after them and could go back to their computers, programs, and networks.

Although Stuttgart's simulation model could be communicated to its own user community in the form of programming instructions, communicating its simulation competence to the outside world was done visually. The Stuttgart "tetrahedron" helped in this effort by abstracting the individual processing steps during a visual simulation and representing them schematically as basic functions of problem solving (Fig. 7). The representation enabled a unified view of the multifaceted process of visual simulation at RUS. It was now possible to show competitors, sponsors, and observers that by combining "calculation, representation, supercomputing, and visualization" in an "integrated system," simulation could be a tool for research and development of theoretical models in engineering sciences.[82] This system enabled the interdisciplinarity that was the dream of computer scientists: the collaboration of different disciplines under the aegis of those who understood the computer better than the scientific problem and who themselves also worked in an interdisciplinary field of knowledge. According to one of the the university's research reports, high-performance computing problems, which had defied solution for several years, required "collaboration between specialists from the application areas of numerical mathematics, computer science, and computer technology."[83] This twofold interdisciplinary orientation was programmatically generalized through the invention of computational science: "Instead of research in individual fields, systems must be integrated in all their complexity. The result is a new discipline, which the Americans call 'computational science.'"[84] RUS's own competence in the field of visual simulation, understood as a method of computational science, acquired accessible communicative news value in the community and among funders.

The false security of benchmarks

Since Stuttgart's pivot to the future, with Lothar Späth's purchase of the world's fastest machine, RUS had embraced a supercomputing future – a future that seemed reliably predictable (Fig. 8). The speed and performance of top-flight computers would continue to increase steadily, especially since parallel computer systems were now commercially viable.[85] Nor would the demand for high-end computing capacity diminish. There was "agreement within the community that there would continue to be an almost unlimited need for computing in science and industry in the future, requiring ever more powerful supercomputers."[86]

It was therefore hardly surprising that in purchasing the Cray-2, the state of Baden-Württemberg had already secured an option on "the first Cray-3 to be delivered to Europe, and certainly the first to be installed in Germany."[87] To ensure its contractual autonomy, Stuttgart submitted to the DFG in 1989 a "pre-proposal" for procurement of a Cray-3 – which had just been announced by the manufacturer – in 1991. The anticipated Cray-3 seemed real enough and would be much more than a mere successor to the Cray-2. For the limit of computability would be raised dramatically. And it would be needed. Problems that required more than 100 hours of CPU time on the Cray-2 were hard to solve. With such "top simulations," Rühle had told *Computer Magazin*, problems were nearly incomputable, and simulations "simply no longer [made] sense."[88] With the Cray-3, Cray Research promised to replace the silicon chip technology, thus overcoming the current physical limits of computers. In 1989, Cray's German representative, Robert Übelmesser, announced to the audience at the Mannheim Supercomputer Seminar that "under the leadership of Seymour Cray, the Cray-3 project" was developing a supercomputer using "gallium arsenide technology."[89] Owing to "faster switching

zurück

times and lower heat dissipation," the technology offered "the potential for even more powerful supercomputers."[90]

The Cray-3 had long since become an integral part of Stuttgart's planning, and was included in the proposal to the DFG. But focusing exclusively on a future high-performance system would likely appear too risky to the DFG, which was more inclined toward broad-based research. Instead of relying solely on the performance of the machine, the Stuttgart proposal also demonstrated the compelling performance of RUS as a whole. The DFG's assessment of need should not be based on the usual supercomputing criteria. Neither the theoretically possible speed of the computer nor its price–performance ratio should be decisive in funding the Cray-3.[91] Rather, the DFG's decision to bankroll a Cray-3 for the University of Stuttgart would be a judgment on the profile of Stuttgart's simulation culture. Accordingly, the "size and experience of the user community,"[92] the on-site software resources, and the bandwidth for accessing the computer from the network and workstations were highlighted.[93] The increased processor performance and memory, and the faster memory access time expected of the Cray-3, were additional selling points. The performance enhancements of the future computer would only be effective if all the criteria were "balanced."

The political strategy behind the pre-proposal was to use state policy to apply pressure to the DFG. It should appear to the reviewers that Baden-Württemberg had already factored in a positive decision from Bonn. The state was awaiting the DFG's decision "so that the Council of Ministers and Parliament [could] proceed with the corresponding budgetary decisions."[94] "Delaying the procurement procedure" would simply be "irresponsible" in view of the "time-consuming approval process."[95]

What might at first appear to be greater self-confidence on the part of Stuttgart vis-à-vis the DFG arose out of a dilemma that

zurück

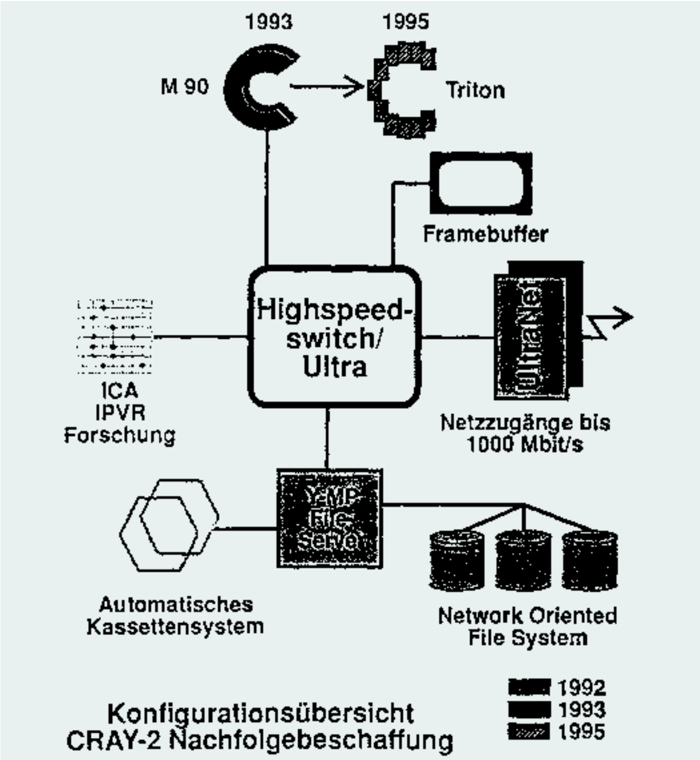

Fig. 8: Stable expectations of future changes, 1992: a configuration overview.

was deferred to the summary at the end of the proposal as a precaution, to wit: the "temporal planning" of the procurement of a successor for the Cray-2 was "of crucial importance."[96] The most decisive factor, however, was competitive pressure on the Stuttgart model of supercomputing. It was no longer certain that the existing set of determinants for operating a regional supercomputer would continue to hold in the future. "Deciding too early on a product type," the DFG's reviewers learned while reading the proposal, "is just as risky as failing to choose a successor computer in a timely fashion, which could mean loss of a leading position in global scientific competitiveness with all the consequences that entails."[97]

The mention of risk evoked a global world of high-performance computing dominated by competition and regional economic policy. In Stuttgart, heightened uncertainty was expected to persist into the 1990s. The question of the right product type and, consequently, the most suitable supercomputer architecture was difficult to assess. It was unclear whether it would be possible to keep up in the global race for the most powerful machines, or whether there was even a threat of losing the pole position.

The market for high-performance computing gained momentum in the early 1990s and became confusing for observers of the field. The "Formula 1 class"[98] of supercomputer manufacturers saw new entrants from Japan – Hitachi, Fujitsu, and NEC – which increased competition between suppliers. The NEC SX3, for example, as Hans-Werner Meuer reported from the Mannheim Supercomputer Seminar in 1991, had "quickly gained a foothold in Europe."[99] Such a machine was ready to be installed, for instance, in Ticino "as the national supercomputer of Switzerland."[100] Parallel architectures increasingly provided a viable alternative to the dominant vector computers for computing. Equipped with many low-cost microprocessors and distributed memory, parallel computing promised to boost performance based on an improved price–performance ratio.[101] But this form of computing, as with other innovations before it, required a different way of programming, for which neither the appropriate software nor expertise existed, at least not in Germany.[102] Workstation clusters offered an attractive, cost-effective interim solution, so long as it remained unclear "which of the many proposed concepts" for parallel computer architectures would "stand the test of time and prevail."[103] Computer development in the early 1990s was therefore considered to be "essentially unpredictable over a period of more than 3–5 years."[104]

In the field of high-performance computing, the evaluation issue was exacerbated because machines like the Cray-3 did not even exist.[105] Could the manufacturers' promises of the theoretical maximum performance of their new machines be believed? Was price an indicator of the actual performance of the computer? Evaluating the next generation of supercomputers remained a difficult-to-solve task for the personnel responsible for new purchases in the computing centers. In Stuttgart, the situation was such that "critical manufacturer information on certain data was actually not available and quite possibly not available for the foreseeable future." How was it possible to properly justify the choice of this or that supercomputer in such a situation? The only options left to computer centers were to take it upon themselves to painstakingly extrapolate the performance values of existing computer systems, to ask their colleagues, or to leverage good relations with vendors.[106]

But extrapolating the possible performance of a future computer system based on the simulated performance of a known computer jeopardized the credibility of everyone involved. Even comparing parallel and vector computing could bring accusations of scientific inaccuracy from the applied mathematics community.[107] Scientific staff in computing centers were constantly being told that reference conditions could not be reproduced or, to the contrary, that they could be readily adapted for the advantage of the center's own computers. It therefore quite frequently happened that the results of performance comparisons between the same vector computers differed by a factor of 10. A paper published at the time observed that "performance measurements on factor and parallel computers are very difficult to compare with universal computers."[108]

The pressure on scientific staff could hardly have been higher. In submitting its confidential internal study on the performance of four high-performance computers, the numerics for supercomputers group at RUS observed that the laboriously evaluated benchmarks were basically worthless: "The data obtained using a simplified model for vector processor circuits cannot be claimed to be accurate."[109] The study's authors concluded that "it was entirely possible that essential architectural features of the machines were not taken into account and, consequently, that the predictions of performance were incorrect."[110]

A solution to the confusion in the manufacturers market for supercomputers – the Top500 ranking – was published on 24 June 1993 at the eighth Mannheim Supercomputer Seminar. Taking its cue from *Forbes* magazine, which had listed the 400 richest Americans since 1982, the Top500 ranking was intended to provide information about the world's most powerful computers.[111] Not that manufacturers and users had previously been flying blind. A number of existing directories provided overviews and mapped developments. Because many universities in Europe and the United States were equipped with powerful machines, national funding agencies, for example, kept site lists. Supercomputer locations were listed in alphabetical order, along with the number and types of machines, and frequently their specific purpose.[112] Or the lists might specify whether the computers were located in the computing center, the library, the university administration, or the local institutes.[113] It was even possible to create lists of the number of vector systems worldwide, as Hans-Werner Meuer and Erich Strohmeier had been doing at the University of Mannheim since 1986.[114]

The Top500 ranking marked a fundamental departure in the universe of surveys. The product of a collaboration between Meuer and Strohmeier and the American mathematician Jack Dongarra, the ranking simplified surveys dramatically.[115] It did away with manufacturers' specifications, the market price of the computer, and the many individual computing center benchmarks. Evaluation of supercomputing was reduced to a single criterion: performance. It was measured using computing tests, in this case, the LINPACK Benchmark. The LINPACK algorithm "required the computer to solve a tricky system of equations that just fit into the working memory."[116]

By combining benchmark and ranking, the Top500 list established a reality based on competition. If, according to the LINPACK Benchmark, performance was the only relevant criterion for comparison, then it was no longer the locations that determined the supercomputers. Rather, it was the other way around. The performance order of the fastest computers defined the worth and importance of the sites. In keeping the number of computers on the list to 500, the ranking also functioned as a tool of exclusivity. Mainframe computers, workstations, and "mini-supercomputers" were now banished from the club of top computers. Organizations that wished to compute at the cutting edge with older or cluster-type models were also left out. Whereas those who made it onto the list –preferably the very top – belonged to an elite that could join the race at the limit of computability. But the positions were fluid. Because the ranking was published twice a year – in June at the (newly titled) International Supercomputing Conference in Mannheim and in November at the IEEE Supercomputing Conference in the United States[117] – the top positions could change at intervals of only a few months, new computers and institutions could be added, and others dropped out. The Top500 list provided a measure for determining the excellence of a country's own research centers and, at the same time, for declaring global competition with other locations.

When the first Mannheim ranking appeared in June 1993, the University of Stuttgart's supercomputers were way down on the list: the Cray-2/4-256 was in 205th place, and the MP-1216 was ranked 430th.[118] Stuttgart repeatedly stressed that "in supercomputing, the theoretical maximum performance is completely irrelevant and the LINPACK Benchmark data nearly as irrelevant."[119] Nonetheless, the close cooperation "with the American NSF supercomputing centers, the Los Alamos National Laboratory, Livermore, and NASA Ames, as well as the Cray and Convex computer companies"[120] did nothing to change Stuttgart's poor showing. Stuttgart's existing strategy for running supercomputers was not globally competitive. Moreover, long-standing supplier Cray also proved unreliable. Four years after the "pre-proposal" for the successor to the Cray-2, and seven years after its installation, the Cray-2 was still in operation in Stuttgart.[121]

Autonomy through regional cooperation

The Stuttgart model of offering supercomputing services through a university computing center was under pressure from the Top500 ranking. A new, state-of-the-art high-performance computer was nowhere in sight. The dynamics of the manufacturers market due to the new Japanese entrants, and the avidly debated question of whether workstation clusters might not be an inexpensive alternative to high-end computers, heightened the risk of any decision.[122] In Stuttgart, an attempt was made to offset the pressure by reconfiguring the organizational structure and service offerings of the computing center, calling on the store of experience of RUS and the university administration.

Since the acquisition of the Cray-2, high-performance computer operation at RUS had taken a variety of organizational forms, realigning the center's relations with the periphery. Along with the installation of the supercomputer in 1986, a scientific and technical computing group was formed, and the new management position of scientific director was created. The first director was Roland Rühle, who also held a chaired professorship in applied computer science in mechanical engineering at the University of Stuttgart. The idea of the dual appointment to establish "a connection between the computing center and research and teaching"[123] at the university – in other words, to find new local users for the supercomputer's capacity. Three years later, in 1989, the aim had morphed into motivating the government of Baden-Württemberg to initiate "timely procurement" of the successor computer to the Cray-2. Accordingly, RUS pledged to establish a "center for the simulation of science and engineering systems."[124] In the end, there was no desire in Stuttgart to follow the recommendation of the "future committee" to build a stable connection as far as Karlsruhe. As a result, all talk of a joint center for simulation ceased.[125]

Instead, new links were established in Stuttgart between the university and RUS. In 1993, the Competence Centre for High Performance Computing was unveiled to the public.[126] To this end, connections were forged between the Institute for Computer Applications II (ICA II), the Institute for Parallel and Distributed High Performance Computers (IPVR) and RUS.[127] The purpose of the center was the "interdisciplinary development of supercomputing and its applications in research and industry."[128]

Through these connections to the University of Stuttgart, RUS was able to acquire competence in parallel computing, gain the status of an independent research unit for funding programs, and intensify its working relationships with the regional automotive industry. Cooperation within the competence center was expected to lead to greater autonomy for supercomputing in Stuttgart.

The collaboration between the two university institutes now made it possible to experiment with massively parallel high-performance computing in the computing center. RUS had gained experience through the "modest parallelism"[129] of the Cyber 205 and Cray-1. In the Cray-2, RUS possessed a machine with several processors that could execute jobs independently of each other or work simultaneously on the same program.[130] Nor was the concept of massively parallel computer architectures unfamiliar: the famous Illiac IV in Illinois had been around since the 1970s.[131] "It's an old idea,"[132] asserted Alfred Geiger, head of the department for parallel computing at RUS, before explaining to his user group – all used to vector computers – why "the RUS, IPVR, and ICA II" wished to "jointly install an Intel Paragon XP/S-5"[133] in December 1992.

According to Geiger, the advantages of parallel computing "[were] obvious."[134] Because parallel systems could be expanded at will using standard components, and because the scalability of the systems "and thus the potential performance [are] essentially unlimited," these supercomputers offered "significantly greater gross capacity ... for the same price."[135] In theory, parallel computers were faster than vector computers. In practice, however, they required careful, time-consuming preparation.[136] Jobs first had to be parallelized, otherwise the system performance of the computers was significantly below par. Compensating the difference between practice and theory meant that users first had to identify parallelisms in their complex problems and then translate them into algorithms. For this step, vector computers offered tried-and-tested, standardized compilers that supported users in translating their problems. The special requirements for different types of parallel computers – whether workstation clusters or supercomputers – as well as the different programming models thus led to considerable "effort on the software side."[137] Programming massively parallel supercomputers was more complex than the well-rehearsed routines of vector computers, leading RUS director Rühle to conclude: "Programming efficient software in the science and technology area is getting more difficult, not easier."[138] Making productive use of parallelism depended on the "efforts of the user."[139] That, with respect to the computing center's operations, was the surprising result of the new computer architecture: parallel computing meant the return of the user.

With massively parallel computing, users had to leave the workstations assigned to them by RUS's integrative operating concept and return to the computing center. Once again, they had to be aware of which supercomputer would be calculating their jobs. Parallel supercomputing was not something that could be done at a distance. Users needed to be close to the machines to be able to intervene in the individual processing steps to minimize latencies. Interaction between the parallel computers and their potential user community therefore had to be carefully defined in the computing center to keep the existing operating concept. This appeared to be much easier than to conceptually rework the well-established, highly sophisticated procedure.

Responsibility for this task fell to the competence center. Because user support in massively parallel computing "would by far exceed the scope of guidance,"[140] new organizational forms for operations had to be sought. Here, too, the center was to be strengthened through cooperation, and it was decided that operations in parallel computer architectures should "take the form of project partnerships."[141] The division of labor vis-à-vis use of the machines and the "development work"[142] on problems of parallel computing was organized in the form of projects between the computing center and the university.

Like the Intel Paragon machine, the IBM RS/6000 PARIS cluster was also acquired in 1992 as a project. The project on workstation parallelism involved RUS; Collaborative Research Center (*Sonder-*

forschungsbereich) 259 on "High-Temperature Problems of Reusable Space Vehicles," in which the mechanical and production engineering departments participated; the German Aerospace Center; and IBM. An additional parallel computer, the Cray T3D, was purchased in 1995 in cooperation with RUS, ICA III, and Cray Research.[143]

The cooperative strategy was also strengthened symbolically (Fig. 9). In an issue of *BI. Informationen für Nutzer des Rechenzentrums*, a newsletter for the computing center's users, Alfred Geiger provided users with an introduction, complete with pictures, to the center's staff, "who will strive to support you in solving your problems."[144] Not every staff member was pictured, but readers nonetheless learned that Manuela Sang, for example, was responsible for "user support (parallelization, libraries, and applications) on the PARIS cluster" and "also for all questions related to the PVM and FORGE 90 tools."[145] Readers were also introduced to Manuela Zürn of the numerics for supercomputers group, who was working on a project funded by the European Community for porting computational fluid dynamics (CFD) code to parallel computers. Finally, readers learned that Gabriele Schulz-Ziemer was working on a project "in which, for the first time, the HPF programming model"[146] was being used on a parallel computer to solve a practical problem.[147] It took many RUS staff to ensure user autonomy in massively parallel computing.

RUS now also acquired a past in parallel computing. Through the connection with the institutes, it could be claimed that massive parallel computing was already taking place in Stuttgart "before there was any name for this type of processing."[148] Although a parallel computer had only been available in the computing center since 1992, preoccupation with problems of parallelism stretched back "far into the previous decade."[149] RUS research projects also integrated new users inside and outside the university into the computing center. For example, the computer simulation and vi-

Fig. 9: RUS staff, 1995: How can we help you?

sualization department at Rühle's Institute for Computer Applications was working on the "integration of calculation programs" in a "joint project with Daimler Benz AG."[150] Thus, alongside Porsche, tentative ties were also being formed with the region's largest automobile manufacturer. Three engineers from the applied mathematics in mechanical engineering degree program worked on the project as contractors for Daimler-Benz, as the rector proudly noted in the annual report.[151]

Integrating the computing center into the university's various institutes had turned RUS into an autonomous research institution. Thanks to the EU's PAGEIN (Pilot Applications in a Gigabit European Integrated Network) project (1992–1995), the French Aerospace Lab (ONERA) became a customer, as did the German Aerospace Airbus company through the EU's ADONNIS (A Demonstration Of New Networking Integrated Services) project (1994–1996). During the same period, RUS was also involved in several collaborative research centers and in projects related to state research foci.[152] As a byproduct of these various cooperative efforts, RUS had access to new external sources of funding from

European and national funding programs, making it possible to "pursue strategic development that would not have been possible using [its] own resources."[153]

Stuttgart's two-pronged solution

Networking and cooperation were the definitive forms of coordination for supercomputing operations in Stuttgart in the early 1990s. In the competence center, projects, experience, and expertise in massively parallel computing and access to high-performance computing capacity were mutually shared between the university's institutes, the local automotive industry, and RUS. This network of relationships enabled the computing center to reduce the uncertainties surrounding developments in computer architectures and the programming of massively parallel algorithms and software. But even with the competence center, there was still no successor machine for the Cray-2 to offer in Stuttgart. The obvious question was: Couldn't the cooperation and networking model also be leveraged to finance a future high-performance computer that would end up on the Top500 list?

The DFG's committee on computer systems had in principle "deemed worthwhile Stuttgart's proposal to procure a high-end supercomputer."[154] However, both the DFG and the Science Council were of the opinion "that such a computer should be accessible not only to Baden-Württemberg but also to universities throughout Germany."[155] High-performance computing, as the DFG's suggestion made clear, had become an attractive area of action for federal policy. The computing centers, which operated high-performance supercomputers and were always looking for viable funding models to procure their expensive machines, saw the new interest on the part of policymakers as an opportunity.

It fueled the transformation of the supercomputer into a demonstration object of national competitiveness.

In early 1992, the High Performance Scientific Computing Initiative (HPSC), which brought together the who's who of German supercomputing from Hamburg to Munich, had outlined a series of recommendations for policymakers in a position paper. Right at the outset, the paper referred to the "key role" of high-performance computing, and asserted that high-end computers were becoming "decisively important for ensuring the competitiveness of the German economy."[156] The Top500 list provided policymakers with an indicator that was as simple as it was clear and that made competitiveness visible to all. There was no denying that the results of the initial rankings were sobering for the country. In 1994, Germany's high-performance computers achieved "less than 7% of the performance offered by the world's leading computer system."[157] In "international comparison," the country had thus fallen not only behind Japan and the United States "but also behind France, Great Britain, Italy, and Switzerland, as well as Canada and Korea."[158] However, with the purchase of just one new machine (a tempting thought born of the biannual publication of the Top500 list), Germany could reclaim the forefront as early as tomorrow. Accordingly, the federal government was only interested in computing capacity that was "competitive with the highest performance class internationally."[159]

Federal policymakers and the computing centers may have seen eye to eye on the semantics of a national supercomputing system. But the financing and operation of future high-performance computers was not yet secured. The reunification of the two Germanies both tied up attentional resources and increased the number of higher education institutions eligible to apply for funding. Even Baden-Württemberg was facing "difficult economic times" and "tight budgets."[160] Nevertheless, the federal

zurück

government shifted responsibility for financing the Stuttgart supercomputer to the states. This posed a real problem for German federalism. On the one hand, universities throughout the country were supposed to have access to Stuttgart's top computer. On the other hand, it was up to the state of Baden-Württemberg to finance the machine. This conundrum gave rise to the question of how "participation by all the federal states in a joint computer in Stuttgart could be achieved."[161]

A proposal by the Science Council to establish "centers for the supra-regional supply of science and research with high-performance computing capacity"[162] was of little help in this situation. The council's suggestion that states and universities should examine "whether planning should be carried out in a network of several states" was hardly persuasive,[163] given that the same paper called for "competition between locations."[164] Although the federal government made "additional project funding" available outside of the University Construction Act in view of the "special technological importance" of supercomputing – in other words, proposals for a national supercomputing center "[did] not count toward the proposal quota for the respective state under the [Act]"[165] – the question of the percentage of financing to be supplied by the "home state" remained open. The Gordian knot of national high-performance computing would have to be cut locally, in Baden-Württemberg and in Stuttgart.

The model of cooperation and networking for high-performance computing operations that had been tested in Stuttgart was attractive for this purpose. It promised to take the various interests into account and to regulate the required participation of the federal, state, and university authorities. For the University of Stuttgart to remain a competitive location for supercomputing, the local model had to be scaled to Germany's policy situation. To maintain its foothold, Stuttgart would have to reaffirm its own identity. What was the mission of RUS? What had it achieved? And what should be its stance toward the future?

These questions were answered in a brochure from 1994, which bore the simple title *Das Rechenzentrum der Universität Stuttgart* (University of Stuttgart Computing Center). The very first sentence of the brochure connected the present and past, assuring a solid foundation for authors and readers alike. "The goal of the Regional Computing Center of the University of Stuttgart (RUS) is and always has been to provide its customers with the most powerful systems for scientific and technical computing."[166]

The most powerful machines were then introduced in chronological order: "The major milestones are the CDC-6600 from Control Data, CRAY1/M, CRAY-2 up to today's C94D with 8 Gbytes of main memory."[167] From there, the text turned from describing the center's own supercomputing supremacy to emphasizing its mission to provide basic regional services, the proximity and collaboration with users, and local problem-solving competence. At issue was the "operational security, service quality, and availability," the export of computing time "also offered to users outside Stuttgart," the network infrastructure, and finally "data management" in the context of decentralized computing.

RUS's basic service mandate was also expressed in a simple organizational diagram that depicted the departments, areas of responsibility, and staff functions at the computing center. Roland Rühle was at the top of the chart; the base consisted of three departments – "High-Performance Computing," "Service and Networks," and "Basic Services." Readers could see that RUS had much more to offer than just supercomputing: "We hope that this organization will enable us to meet the challenges of the coming years."[168] With this sentence, the introduction then segued to detailed enumerations of the services and current research projects at RUS.

Rühle's hopes were disappointed. Scaling the cooperation model to the federal level disrupted the existing organizational structure of RUS. Up to that point, networking had previously been supported by personal contacts between the parties involved or by sharing staff between the computing center and the university institutes. Now, the focus on national supercomputing meant a considerable expansion of cooperation to the federal and state level and to industry. The new model required legally, politically, and economically strong connections. These could no longer be achieved through the prevailing organizational forms, which were geared to high-performance computing operations at the regional level.

National supercomputing in Stuttgart began with the establishment of two new organizations.[169] In 1995, the Höchstleistungsrechner für Wissenschaft und Wirtschaft – HWW GmbH (hww), a public-private partnership – was founded, followed by the High Performance Computing Center Stuttgart (HLRS) in 1996. A principle of separating capital investment and allocation of computer time shaped national high-performance computing in Stuttgart. The task of hww was to secure financing for future supercomputers. The HLRS, for its part, was responsible for allocating computer time to German universities and research institutions.

What appeared ostensibly to be a straightforward division of organizational labor was actually highly contentious. In her address at the inaugural ceremony of the HLRS, University of Stuttgart rector Heide Ziegler summed up the challenge: "Behind us lies a phase of several years of planning and preparation, discussion, and at times also disagreement."[170] When the HLRS was founded, it was still possible to fall back on the established networks of cooperation at the University of Stuttgart, although they seemed to have become somewhat looser with time. Ziegler suggested as much in her reference to the "protracted yet ultimately

zurück

successful interdisciplinary and inter-institutional effort" necessary to establish the HLRS.[171] The founding of hww was a different story. What the state government would later laud as an exemplary model of cooperation between science and business, for a long time plausibly met most of the criteria for communicative failure. In her annual report, Ziegler had carefully referred to hww as a company with a "somewhat complicated (and laboriously negotiated) name."[172] By the time of the inaugural address, however, her phrasing was more pointed. Some in the audience no doubt looked at their shoes as she explicitly thanked those "who were involved in the long and, due to the diversity of cultures, sometimes almost insoluble task of creating hww GmbH."[173]

Elsewhere, the new Stuttgart organizational model for operating a national high-performance computing center was viewed more euphorically. Observers from the German high-performance computing community interpreted hww's founding as a "Stuttgart coup."[174] Rumored to have been "engineered at the very top" by "Edzard Reuter (Daimler) and Erwin Teufel (Baden-Württemberg),"[175] hww was a model demonstration of modern capitalist cooperation between science and industry. The company was based on a shareholder model. The state of Baden-Württemberg, the University of Stuttgart, the Daimler subsidiary, IT services provider debis Systemhaus, and Porsche participated as shareholders to varying degrees.[176] The cooperation specified two conditions. First, participation should be open to other universities. Within the state government of Baden-Württemberg, the University of Karlsruhe was high on the agenda.[177] Second, parity of participation was to be maintained between industry and science, with each side holding a maximum of 50 percent of the share capital. The chairman of the shareholders' advisory board, appointed by the state of Baden-Württemberg, supervised these rules of participation.[178]

The new HLRS was organized in a completely different way, namely, according to Germany's system of proportional representation.[179] The interests of the federal and state governments, the other universities in Baden-Württemberg, and the University of Stuttgart all had to be carefully balanced in terms of allocation of computing time. The exclusivity of the offer also had to be guaranteed. Operators had to ensure "that the new high-performance computers were used exclusively for major problems that could not otherwise be executed on any computer."[180] The HLRS main body for this purpose was the steering committee. It consisted of twelve members; six were appointed by the DFG, and six by the Baden-Württemberg rectors' conference.[181] The joint committee would establish rules for the allocation of computing capacity, decide on project proposals for allocation, and have a say in the selection of hardware and software.[182] The steering committee was also to ensure the separation of investment and computing time and the allocation of fixed usage quotas: Nearly half of the total capacity was reserved for users throughout Germany; approximately 30 percent was reserved for universities in Baden-Württemberg; and roughly 20 percent for local needs at the University of Stuttgart. Of the capacity reserved for users throughout Germany, only 8 percent was available to industry.[183]

The cooperative arrangement between hww and the HLRS created a national high-performance computing center toward which all parties involved – the federal government, the state of Baden-Württemberg, the University of Stuttgart, and industry – could focus their activities. The center strengthened the government's claim on Germany's competitiveness in global high-performance computing. The essential differences between hww and the HLRS testified to the cross-system cooperation between the federal and state governments and between industry and science. Through its "supercenter,"[184] the University of Stuttgart acquired an exclusive facility for computing at the limit of computability.

The basic task of supplying the university with computing capacity was taken over by the new University Computing Center department at RUS. Local users, who had enjoyed close collaboration on numerous projects until recently, were also affected by the new exclusivity of high-performance computing in Stuttgart. They were now customers like any others in the German high-performance computing sector. Henceforth, the department of adaptive structures in aerospace at the University of Stuttgart had to apply to the new steering committee of the HLRS for its allotment of computing time, as did the department of mechanical process engineering at the University of Kaiserslautern or the aircraft systems engineering department at the Technical University of Hamburg-Harburg.

The new high-performance computing center, on the other hand, functioned as a virtual center. It depended on other organizations for its "operability."[185] According to the operating concept, the HLRS had no computers of its own, but rather purchased computing time from the operating company and passed it on to scientific users under the terms set by the steering committee.[186] Responsibilities were also decentralized. The HLRS took over the office of the steering committee, but the committee itself was organizationally linked to the Baden-Württemberg ministry for science. The new center was also linked with other German centers of excellence in the area of high-performance computing with the aim of achieving "participation of all interested parties in the work of the high-performance computing center."[187] The connection with the center's local base was also redefined. "In order to meet the requirements of the federal high-performance computing center in Stuttgart," stated Ziegler's rector's report succinctly, "RUS has created a high-performance computing department and a university computing center department."[188] The national computing center had freed Stuttgart supercomputing from all its former ties, made it accessible nationwide, and distributed responsibility for its operation throughout the organization.

The intricate model of cooperation between hww and the HLRS was built with the ideal of setting an example. The hope was that it would radiate to other centers as an exemplar of globally competitive high-performance computer operations. Certainly no event radiated more brightly than the state press conference of 30 January 1996, when government and business leaders in Baden-Württemberg gathered in the cramped quarters of the University of Stuttgart's computing center. This group was joined by Erwin Teufel, the state's minister-president, and his science minister, Klaus von Trotha. The first row on the podium was occupied by captains of industry: Gerhard Barth, head of information technology corporate research at Daimler-Benz; Walter Gnauert, chief financial officer of Porsche; and Wilfried Steuer, chairman of the supervisory board of Energie-Versorgung Schwaben AG (EVS), who was also introduced as the head of "Communikationsnetze Südwest GmbH," a company that had just been founded by the regional energy suppliers.[189]

A few bits of news were shared for the benefit of the many media representatives, including the concept of the national computing center and the investment by Communikationsnetze Südwest GmbH of 42 million marks in the high-speed network.[190] However, the focus of the press conference was the announcement of additional cross-system cooperation. With the help of hww, contracts were being signed with Cray for a new parallel computer and with Japan's NEC for a new vector computer. After many years, the succession of Stuttgart's Cray-2 was finally assured. This Stuttgart double coup was elaborately staged. "Following initial questions directed to the minister-president on current state and federal issues by the journalists in the room, which is usual for a state press conference," noted the RUS user bulletin, "ISDN video conferences to Cray Research in Minneapolis and to NEC in Tokyo were activated."[191] Positioned between Teufel and the science minister, Heide Ziegler greeted Phil J. Samper (Cray) and Masao Toka (NEC) – displayed on two monitors and one projection screen – with a warm "Welcome to Stuttgart!". At the end of the press conference, Stuttgart's strengths and expertise were once again put on show. Participants were able "to watch on a stereo-projection screen a three-dimensional representation of an aircraft body that changed color according to the distribution of pressure on the surface."[192] The research partners involved in the project could be seen on the screen via the simultaneous video conferencing.[193]

The showcasing of the Stuttgart cooperation model for national supercomputing was by no means limited to demonstrations of technical feasibility at a state press conference. The "model character"[194] of the Stuttgart cooperative venture was also intended to apply beyond national high-performance computing. On the occasion of the HLRS inauguration on 12 September 1996, science minister von Trotha mused aloud about the new center as a model for other areas of science policy. The HLRS, he said, struck him "as a good example of future-oriented university policy, since here the tasks of the university are privatized in a commendable manner."[195] Von Trotha also recognized the importance of the HLRS "having no computers of its own" and "in this respect being a 'virtual computing center.'"[196] The new national supercomputing operation in Stuttgart had become a viable model for Baden-Württemberg policy in the 1990s: "Just like the users in the computer center, the student of the future will less frequently be encountered at the university."[197]

Network to the rescue (1997–2005)

Essentially everyone who predicted a bright future for the HLRS was at the state press conference in January 1996: the minister-president, the rector, representatives of the auto industry, transmission line builders from the electricity industry, the long-time director of RUS, and many others. One manager each from NEC and Cray were beamed in from a sleeping Tokyo and pre-dawn Minneapolis. Nobody noticed that the federal government was accorded no speaking time. Whether the omission reflected a wish to keep the protocol simple or because the national component of the HLRS had not yet been finalized is difficult to say.[198]

There may also have been a reluctance to overdo things. After all, the impressive title Federal High Performance Computing Center had been secured, as had the purchase of new machines. And industry had finally been persuaded to create hww. After endless hours of meetings, strategy papers, presentations, telephone calls, reports, proposals, and interviews, Roland Rühle could be proud of having safeguarded high-performance computing services in Stuttgart for the long term. The process would continue. Nonetheless, the combination of the HLRS and hww now opened up the possibility of flexible accounting in the collaborative arrangement between university and industry. Moreover, the near simultaneous procurement of the NEC SX4 vector computer and the Cray-3TE system for massively parallel processing enabled Stuttgart to offer both conventional and experimental high-performance computing services. Stuttgart had really reached the top, even of the Top500 ranking.[199]

zurück

It was no secret that the federal funds would be insufficient to cover several of these systems every fiscal year. It was also clear that the amount of computing to be done in Stuttgart was less than the capacity made possible by the machines. What was installed on behalf of the federal government in Stuttgart, and later in Munich and Jülich, led to unprecedented computational overcapacity. One could only hope that institutes in Hamburg, Aachen, Bochum, Hanover, Cologne, Bonn, Karlsruhe, and Kiel, which had all come away empty-handed, would help to utilize the new machines to the full. Fortunately, new collaborative research centers and priority programs sponsored by the DFG had been established in Chemnitz, Dresden, Erlangen, Karlsruhe, Munich, and Saarbrücken, which could also be expected to help reduce the overcapacity. However, doing that would require expanding existing connection capabilities.

This need proved convenient for Federal Ministry of Education and Research (BMBF). The recently reconfigured authority began to carve out a highly visible area of science policy in high-performance computing that could potentially also encompass something along the lines of a high-speed network. In any case, in the mid-1990s no infrastructure was easier to justify. It was simply considered timely to advocate “seamless connections between networks,” to promise the competition-neutral safeguarding of “interoperability,”[200] and to enthuse about the great “data highway” that would connect everything with everything else.[201] It stood to reason, then, that rather than having to deal with the niceties of state press conference diplomacy and possibly getting into trouble on behalf of the federal government, the BMBF preferred to deal with the question of how nationwide access to a small number of federal high-performance computing centers could be secured in the near future and how the postal monopoly, which hampered competition, could be circumvented at the European level.

The guiding principle for the work of configuration had thus shifted once again. The problem in supercomputing was no longer establishing centrality or locally orchestrating performance. What had to be relearned in Stuttgart at the end of the 1990s was how to deal with heterogeneous computing clusters. This now involved computers of different providers, various computer architectures, and diverse networks as well as the differences between computing-intensive disciplines and forms of operation at the respective sites nationwide. The focus on heterogeneity in the third configuration model increased the complexity of operations and required a new degree of abstraction. For the willing, heterogeneous high-performance computing in Germany entailed considerable conceptual work for individual operations and building interfaces to other locations.

Feasibility study for a national high-performance computing network

On 18 December 1995 the Society for Mathematics and Data Processing (GMD) made a presentation to the BMBF on establishing and operating a high-performance computing network in Germany.[202] What emerged from the subsequent discussion was a recommendation to (re)study the feasibility of such a network. The GMD had obviously (typically, in the eyes of the federal ministry) proceeded too theoretically and too generally, with little attention to implementation. That, in any event, was Friedel Hoßfeld's reading of the minutes of 3 January 1996. A few months later, Hoßfeld (Forschungszentrum Jülich), set to work, together with Peter Deuflhard and Jürgen Gottschewski (Zuse Institute Berlin), Heinz-Gerd Hegering (Leibniz Supercomputing Centre in Munich), and Roland Rühle (RUS).[203] In October 1997, nearly two years after the GMD's presentation, Hoßfeld and his colleagues annihilated the theorists' work with a new feasibility study.[204]

As a manager of a supercomputing center, Hoßfeld was used to speaking of scientific computing as a "strategic discipline" able to mediate between mathematics and data processing. Such statements were typically bolstered by the oft-repeated assertion that, in addition to theory and experiment, there was a third, complementary pillar of scientific research whose methodology was simulation, whose instrument was the supercomputer, and whose tool was visualization. The almost ritualistic history of science rationale for high-performance computing facilities in Germany also somewhat meanly observed that the importance of these centers had (meanwhile) also been recognized by the Science Council.[205]

In fact, as early as the summer of 1995 the Science Council had produced two separate reports in which it advocated providing science and research with high-performance computing capacity and high-speed data communication.[206] Only in this way would Germany be able to meet the "grand challenges." To this end, the performance pyramid that the DFG wished to see in the German computing landscape also needed a peak.[207] The authors of the feasibility study of 1997 divided this rather broad argumentative framework into smaller justification steps and produced an entire chain of arguments that can be paraphrased as follows: (1) People in Germany realized that high-performance computing was needed, because without supercomputers (2) the country risked becoming a scientific backwater. (3) The German Science Council had also discussed the matter at length and come to the same (obvious) conclusion. (4) The DFG would certainly agree. However, it believed that (5) computers had become routine in science. Therefore, the DFG (6) was not in the business of supporting high-performance computers. The DFG (7) did not even intend to develop an Accelerated Strategic Computing Initiative like the one launched by the US National Science Foundation in cooperation with the US Department of Defense. Unfortunately, (8) this omission had not (yet) been compensated by the BMBF. At least "the federal government" was supporting the Science Council in that (9) it had now finally helped to set up a high-performance computing center in Stuttgart (10) with close links to industry. Moreover, through the operating company hww, a "novel model for closer cooperation between science and industry in the strategic field of scientific high-performance computing" was "actually being tested."[208]

The feasibility study thus adopted the reasoning of the German Science Council, thereby being critical of the DFG, and kept the federal ministry on tap as a possible source of funding for future high-performance scientific computing. Competence structures had been in place for a long time, and "scientific computing" had developed significant "crystallization points." In view of the "de facto existence of high-performance computing centers in line with the structural and capacitive expectations of the Science Council, namely, the supercomputing centers in Jülich and Stuttgart, and in view of plans such as those of the [Leibniz Supercomputing Centre] in Munich and the [Zuse Institute Berlin]," there was now also a "need for clarification and action with regard to overarching strategic perspectives."[209] The rhetoric of competition and rivalry that had evolved in the field of high-performance computing since the mid-1980s and had reached a zenith with the Top500 ranking was clearly off the table again. The federal government was back and would welcome being approached as lender of last resort.

This was a decidedly ingenious tactic. It took the federal government into account all the while relieving it of responsibility. To succeed, the ministry had only to respect the competence in the field, that is, to have the right people do the work. And to provide the financing.

This idea was not original, but it was effective. At the Supercomputer '90 conference in Mannheim, James C. Almond of the University of Texas System Center for High Performance Computing in Austin had already framed his arguments in a similar way.[210] Almond's presentation in Mannheim focused on the "usual financing mechanisms"[211] of state funding agencies. "State planning committees" were used to "financing objects like buildings that last for decades or even centuries. Justification processes that take several years are acceptable in such cases, but not in the case of computer systems that are nearly 40% obsolete in one year!"[212] To Almond's way of thinking, it was impossible to reconcile high-performance computing and the funding policy of the DFG, because the DFG required proof of need for every research and procurement proposal. At this point, he said, there had to be a rethink; policies had to be adjusted. For anyone seeking to calculate at the limit of computability tomorrow and the day after, the one-time purchase of a high-performance computer was no longer enough. "If the university is determined to compete at the forefront of computer-related research," said Almond, "it will need to take a new position. *The mere existence of a new generation of machines must be enough to justify their purchase!* Demand surveys are essentially obsolete!"[213]

Such a challenge was tantamount to disempowering the review boards and relegating the DFG to the role of funder. The logical consequence of Almond's bold thought experiment was that control over funds would then lie exclusively with the computing centers as operators of the supercomputers. Because it could not be assumed "that the state bureaucracy – in Germany, as in Texas – would quickly acquire the necessarily financial agility for such flexible planning," Almond felt that "institutionalizing the computing center as a university body with a secure long-term budget"[214] was the only realistic option. The experienced on-site staff and the sophisticated supply concepts that easily integrated the supercomputers into their heterogeneous machine fleets and network architectures would minimize any danger of "tying the development strategy to a specific hardware and software architecture too far into the future."[215]

Half a decade later, Friedel Hoßfeld and his colleagues would adopt a more sensitive approach to science policy. Strategy was to be left to the policymakers provided the most powerful directors of supercomputing centers could tell the policymakers what the strategy should look like and, in particular, what was "feasible" and should be implemented.

The political situation was conducive to the feasibility study. In any case, the time was right for the BMBF to voluntarily adopt a broad, well-informed strategic view. The ministry had, after all, only been created in 1994 through the merger of the Federal Ministry for Education and Science and Federal Ministry for Research and Technology and was likely amenable to new areas of responsibility.[216]

Hoßfeld and his colleagues were quick to adopt the same big-picture approach. Supercomputing centers should "be able to drive forward methodological innovations in supercomputing and their application in science and research through closer interaction with the other competence centers for scientific computing and their integration into a cooperative network." This would require a "communication network via broadband links" with appropriate interfaces and operating concepts for their "increasingly heterogeneous computer equipment." Only then would "utilization and synergy effects" be possible.[217] However, Hoßfeld and his colleagues had no doubt that the competence for running heterogeneous computer setups lay with the existing large high-performance computing centers, that is, in Stuttgart, Jülich, Berlin, and Munich. As it was, they had all been dealing with difficult concepts for quite some time and, in the case of Stuttgart, even explicitly with heterogeneous computing, also known as metacomputing.[218]

All that was actually needed now was to successfully "pilot test" and expand "B-WiN's [broadband wireless Internet network] broadband communication technology." Accordingly, Hoßfeld and his team could simply have held the federal ministry to account. But working groups, like committees, are diplomatic and do not produce straightforward memos. Rather, after much deliberation, they combine textual building blocks in such a way as to make everything as opaque, exhaustive, and dull as possible. Wolfgang Nagel, who was in charge of editing the feasibility study, would have revised everything and run it by his colleagues several times before producing the following leaden and verbose statement: "In the context of implementing the recommendation of the German Science Council to return German supercomputing centers to the top of the internationally renowned performance pyramid with the simultaneous embedding of competence centers in a cooperative structure for scientific computing, consideration may be given to a network of supercomputer centers with the prospect of fruition."[219]

There you had it. Hard to understand as it was, the text at least showed that the authors understood the political process and knew very well that, unlike their American counterparts, they could not just show up at the ministry with extravagant promises and exorbitant demands. Naturally, they wanted the federal ministry to fund some supercomputing centers. But the authors of the feasibility study also brought something to the table. They had a grasp of their subject, as well as the ins and outs of collaboration throughout Germany. They also knew how scientific computing was done, especially when it came to completely different systems. Without funding from the BMBF, which could quietly also be used to build high-performance networks, Germany (i.e., the German high-performance computing sector) would never find its way back to the front of the pack. That was exactly what the German Science Council, too, had feared, and in the meanwhile had even documented.

The work on the feasibility study, which Roland Rühle took part in, occupied numerous meetings over the space of nearly a year and a half. That was enough time for Stuttgart to review what the feasibility study had determined to be already existing scientific computing competencies, new forms of cooperation, and heterogeneous computing environments. The fact that at the HLRS even a vector computer had to be operationally connected to a massively parallel processor system ensured Stuttgart's familiarity with the problems of heterogeneous high-performance computing. Obviously, Stuttgart had the competence required to deal with a high-performance computing system organized over a network that extended across institutional boundaries.

Metacomputing – a transatlantic experiment

Networking high-performance computing centers could be thought of as an access problem or, similarly to power supply, an issue of load sharing. However, the authors of the feasibility study eschewed such simple analogies. Their mantra would be "Function sharing, not load sharing." In other words, the question was simultaneously one of access and the "diversity of available supercomputer architectures." The functional differences between the machines should be exploited. The wide range of large software systems could also be incorporated into the usage modalities.

All this went far beyond the problem of how to connect a dozen workstations together in such a way as to produce a SCAN – "supercomputer at night" – which had been a focus of study locally since the early 1990s. It also went far beyond the problem of connecting two supercomputers in the same computing center. To this end, the feasibility study developed a bottom-up process to handle a highly complex field. The study was actually comparable with a paper written by the US National Science Foundation and the Department of Defense titled "The Accelerated Strategic Computing Initiative." Because the DFG had not yet come up with a similar program, the study authors had to devise one of their own making.[220]

Fortunately, the first phase of that program, laid out in the feasibility study after a thorough reading of relevant strategy papers and extensive discussion, was essentially complete. It addressed the issue of actual available network capacity. Operations in well-developed high-performance computing centers such as the HLRS could be described, without exaggeration, as a local area metacomputer. Such centers had strong LAN connections and a high degree of integrated hardware and software. They also shared file systems. Indeed, cross-system user coordination created a "seamless environment."[221]

The second phase of the program was more difficult. It was not at all trivial to employ differently designed computers to solve so-called distributed parallel problems. What was to be done with a computational task whose "inherent heterogeneity" meant that "none of the available architectures" could solve it efficiently? How was a field forced by competition into specialization to deal with simulation tasks that required combining different functions?

Solving the problem of heterogeneity in the local computing centers would then (and only then) usher in a third phase that would enable a "(supra)regional or national network from the ensemble of geographically distributed LAN metacomputers via gigabit networks."[222] Such a system was somewhat audaciously called a "B-WAN" or broad-bandwidth wide area network metacomputer. It would be technically highly heterogeneous in composition and could probably handle heterogeneous problems. This prospect was especially attractive "for large engineering applications," which mostly presented as anything but homogeneous. "German supercomputing centers could thus also become international leaders by implementing heterogeneous computing," wrote Hoßfeld and his colleagues.[223]

Massively parallel computing became attractive in the 1990s because the cost structure of computing had changed dramatically.[224] But it wasn't clear how massively parallel systems were to be operated in general. Despite some promising experimental arrangements, fundamental problems arose again and again as soon as a generalized operating concept had to be developed for hardware, algorithms, and software packages that was portable and efficient.[225] Most problems arose from the simple fact that massively parallel computer systems used different solutions for communicating between processors and storage units. There were systems in which computing memory was physically shared, and systems in which the memory was located close to the processors, i.e., widely distributed. Computerized communication could therefore not be standardized because some manufacturers relied on hardware, others on software.[226] Furthermore, it was nearly impossible to transfer existing, well-tested software from a conventional computer architecture to a parallel architecture without having to rewrite everything. Hardly any compilers could automatically parallelize sequential algorithms without affecting program performance. The trio of expectations that programming languages should be expressive, programs efficient, and applications portable could not be satisfied.

When computer specialists cannot simply pass on such problems, but actually have to solve them, either they build a workaround for the locally available machines and their applications, or they circumvent the problem by increasing the level of abstraction and creating a generalized solution. In Stuttgart, a pragmatic combination of both strategies was pursued. First, the specialists drew on their experience with the message-passing interface (MPI), a communication technology standard that had been used as a general-purpose tool for communicating with various parallel computers. MPI allowed operators of high-performance computing centers to protect their customers from heterogeneous machine and application environments. Second, specialists resorted to a connection technology developed in Stuttgart – also tried and tested – called PACX. Both communication technologies had already been combined to establish local connections between different computers: "PACX-MPI (PArallel Computer eXtension) was developed by Thomas Beisel at RUS as a tool to couple an intel Paragon and a CRAY Y-MP," read the information for computing center users under the brief but accurate title PACX-MPI.[227] Its relevance to a high-performance computing network would only be clear once it could be shown that it worked even at long distance. Could Germany's problem of national WAN metacomputing perhaps be solved by taking a detour around the world instead of simply bridging local heterogeneity in Stuttgart? In the midst of the World Wide Web boom of the late 1990s, the idea seemed to suggest itself.[228]

Already in the run-up to the Supercomputing '96 conference, an initial attempt was made to connect the brand-new Cray T3E in Stuttgart with a Cray T3E in Pittsburgh, and to explore their potential interoperability.[229] A further attempt was made in June 1997. Today, the experiment is described by those involved as "heroic but unsuccessful."[230] The heroic part was constructing a virtual supercomputer with 1,024 processors that used a highly heterogeneous connection to exchange data and coordinate applications. This connection led from the supercomputing center in Pittsburgh via NSF's "very high-speed Backbone Network Service (vBNS)" and STAR TAP, to the CANARIE network in Canada, then across the Teleglobe transatlantic network to Deutsche Telekom's network, and from there to Stuttgart – and back again. With the bold announcement that a similar connection to the Intel Paragon at Sandia National Laboratory in Albuquerque would soon be established, the hardware side as well became somewhat heterogeneous.[231] The effort proved worthwhile, however, because Sandia was a strong partner that could contribute a great deal of expertise in operating wide-area, high-speed networks and in visualization technologies.[232]

At Supercomputing '97 in San Jose in November 1997, the metacomputing experts from Stuttgart, Pittsburgh, and Albuquerque presented the very impressive results of several experiments.[233] More followed at Supercomputing '98. In February 1998, the Stuttgart University newspaper, *Unikurier*, reported on three particularly exciting simulations that had been run successfully on a transatlantic network thanks to PACX-MPI.[234]

The first simulation calculated the impact of a comet striking Earth and was worthy of Hollywood (Fig. 10).[235] The initial phase of the catastrophe was computed using a program called URANUS (Upwind Relaxation Algorithm for Nonequilibrium Flows of Stuttgart University), developed at the university's Institute of Space Systems and parallelized at the HLRS. Colleagues at Sandia National Laboratory were responsible for computing the consequences of the comet's impact.[236]

The second experiment was the "molecular dynamics simulation of a crystal of 1,399,400,000 atoms and a granular gas of 1,759,165,695 particles." The University of Stuttgart's Institute

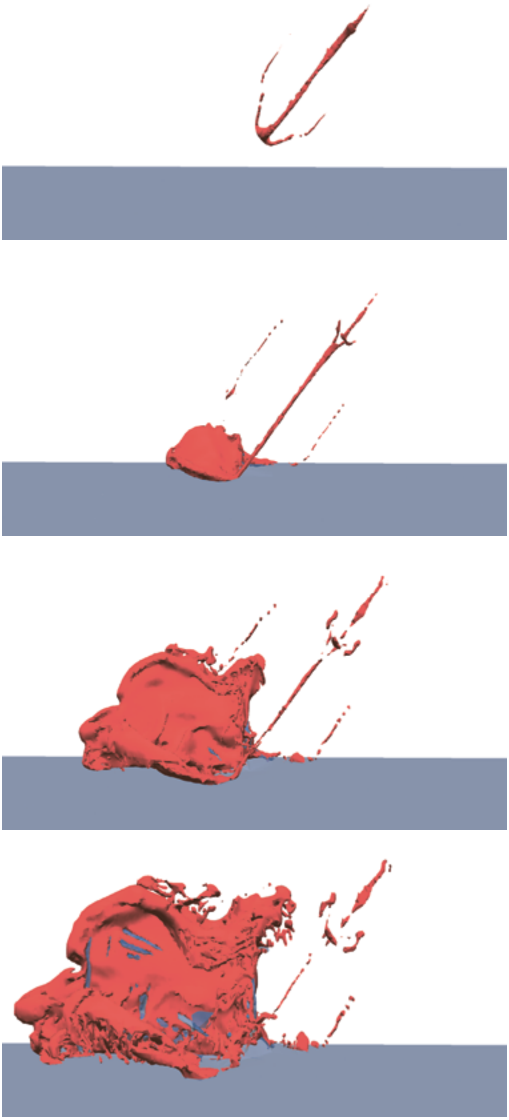

Fig. 10: Metacomputing Armageddon, 1998: simulating a comet strike.

for Computer Applications I had prepared a Monte Carlo particle simulation using object-oriented design. The *Unikurier* described the effort as "world record breaking."[237]

The technology could be taken to a next level if required, namely, in the third experiment demonstrated at Supercomputing '98. This experiment simulated a final stage of the exhaustively planned but never built European space shuttle program HERMES, specifically, a project that NASA had passed on to the European Space Agency involving the X-38 Crew Rescue Vehicle. At the conference, re-entry into Earth's atmosphere after a rescue operation from the International Space Station was computed.[238] This experiment had everything: American, Soviet, and European space travel; American supercomputers; transatlantic cooperation; worldwide networks; and the Baden-Württemberg style of programming with amenability to efficiency, standards, rocket technology, and materials science. Not to mention the huge Collaborative Research Center 259 (High Temperature Problems of Returning Space Transport Systems) with more than 40 subprojects, which the DFG had been funding at the University of Stuttgart since 1990. The report on Supercomputing '98 in the Orlando, Florida-based *HPCwire News Brief* was so euphoric that the successful experiment was even wrongly credited with having demonstrated latency-free metacomputing.[239]

Does the university really need an HLRS?

The final embellishment in the *HPCwire* article was unnecessary. Everything that went before sounded "heroic" enough, at least to European readers. But the mention of latency strained credulity; viewed dispassionately, the big Supercomputing '98 demonstration was even a "failure." Although relatively large bandwidth had

zurück

been secured in the wee hours of European time to cover the long distance from Stuttgart to Pittsburgh and Sandia, the signal time duration repeatedly slowed the painstakingly calculated processing speed. As a result, the metacomputing group could hardly expect a triumphal reception at the HLRS on its return from Florida, Pittsburgh, Albuquerque, and outer space. As was typical of his directorial style, Rühle held everything together. But the increasing skepticism of the current University of Stuttgart rector, Günther Pritschow, was having its own comet-like impact on the HLRS, though less noisily. As an expert in open control systems in flexible production plants, Pritschow was well versed in modern computer use. But his specialty was not known for its affinity for high-performance computing.[240] Rühle was expected to retire in 2003; thereafter, the university would be able to terminate the costly operation of a federally funded high-performance computing center. In the mixed success of Stuttgart metacomputing, the rector recognized the potential for outsourcing high-performance computing. And, as it happened, by the turn of the century outsourcing had moved to the top of everyone's agenda.

It is common practice at universities to see to it that news of such plans spreads by itself. It is not confirmed by declarations of intent, letters, minutes, or articles in the *Unikurier*. At most, it manifests as distilled rumor, for example, in the form of an anodyne announcement of an external expert review.[241] These ad hoc bodies are characterized by their delicate balance of familiarity, over-readiness to criticize, collegial support, and stubborn naivety. The experts are of course given instructions, but they do not wish to be instrumentalized. For every committee is its own self-contained shark tank. Those who leave it alive wish to do so with their head held high and ensure they can return to their own university or company with no loss of face. However, members of a review committee cannot be controlled for the simple reason that in evaluating, they are primarily pursuing their own interests.

Which is not to say that the results of the review process are meaningless. It is simply that the outcome is never what a rector might have hoped, a director feared, or the ministry of education and cultural affairs expected.[242] Moreover, once set in motion, a review cannot be stopped, even when a new rector is more kindly disposed toward the victim of the evaluation. In this case, the new rector, Dieter Fritsch, was a survey engineer who had once written a thesis on the design of digital two-dimensional non-recursive filters.[243] While this topic had nothing to do with supercomputing either, it did have a lot to do with computing and modeling – more, in any event, than Pritschow's communications and control engineering. Fritsch wanted to keep the HLRS at the university, as did Wolfgang Peters, the ministerial representative, who had been alerted by Rühle. Fritsch served as an integrating force and, through careful selection, was able to juggle many balls simultaneously. He may even already have perceived the HLRS as an instrument for his evaluation-based policy of excellence. In any case, the reviewers had to get to work, struggling through thick dossiers and long presentations, highlighting everything that made the HLRS competitive but also criticizing anything that was not satisfactory.

On 15 and 16 May 2000, the external review committee met in Stuttgart. Its composition was prominent, interdisciplinary, and international. The interests of its members remained sufficiently removed from the subject of the review; each member of the committee was aware of his own importance. Rolf Jeltsch, a mathematician from the Swiss Federal Institute of Technology Zurich (ETH Zurich), chaired the committee. Friedel Hoßfeld had traveled in from Jülich. E. Krämer represented DaimlerChrysler Aerospace, and Horst Simon the National Energy Research Scientific Computing Center in Berkeley. W. Thiel was present as theoretical chemist and director of the Max Planck Institute in Mülheim an der Ruhr, and W. von der Linden from the Technical University of Graz represented theoretical, computer-aided physics. The minutes were entrusted to Peter Staub from ETH's Computer Science Services. He was intimately aware of the difficult relationship between a high-performance computing center and a computing center that was slowly diffusing throughout the university, in other words, breaking up.[244]

The Jeltsch committee concentrated on the part of RUS that had to do with supercomputing, that is, only the HLRS. The committee's report showed that even after lengthy PowerPoint presentations, the HLRS structures were not much easier to understand than those of hww. Both were "a complex construct, sometimes difficult for outsiders to understand," but "which function very well in practice"[245] (Fig. 11). The HLRS was embedded in a complex of organizations and was accountable to the management of RUS. However, RUS also managed the internal computing center at the University of Stuttgart (URS). It was clear that the university's interests and the those of the HLRS as the national center for high-performance computing gave rise to conflicts. "At present, these problems are kept at bay only by virtue of the leadership of Prof. R. Rühle," the committee's report noted. And the structures would become "even more complex through the involvement of hww. The state, the universities of Stuttgart, Heidelberg, and Karlsruhe each hold an eighth of the shares in the company. Industrial participation is as follows: Debis (20%), Debis SFR (20%), and Porsche (10%)."[246] Potential gains in synergy would result in uncertainties in the allocation of staff, which again would have to be dealt with by Rühle.

The warning was clear but not particularly grave. External criticism of local complexity is always easier to bear than a lack of

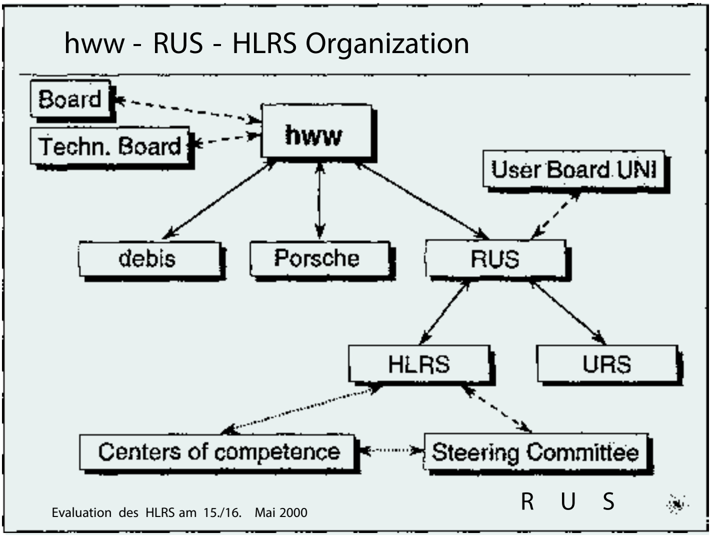

Fig. 11: hww, HLRS, and RUS, 2000: a complex organization that functions amazingly well.

opportunities for local cooperation. At least the review opened up the possibility of intervention by the university administration.

Much graver were the experts' verdicts regarding science. For example, they were surprised by the high proportion of computing time spent on projects devoted to CFD (41 percent), whereas physics accounted for 26 percent and chemistry for 16 percent. There was little demand from electrical engineering, climate research, and biology. The reviewers had no objection to CFD, but somehow the HLRS had to increase its attractiveness for other fields and methods as well. That was the main message of the reviewers' report. It was difficult to say what the post-Rühle era would bring, and they did not wish to hazard any guesses. But because the HLRS needed to widen its appeal to other disciplines, it

zurück

was incumbent upon the reviewers to make specific recommendations. All of them had a special affinity for high-performance computing; together, they came up with a clever proposal based on two mechanisms of integration.

First, scientific computing should be promoted as a course of study. With better educated young scientists, the demand for HLRS services would become more differentiated in disciplinary terms. This was certainly a concern of the committee chairman, who was currently attempting to establish such a course of study at ETH at the interface of computer science, applied mathematics, physics, and theoretical chemistry so that the supercomputing center in Ticino might be able to work on more scientifically interesting projects. However, the recommendation was little more than a rallying cry for a course of study at the university that was unpopular, though no one knew why. The committee was quick to assert that the quality of teaching was not the problem.

Second, there had to be more interaction between the projects at the HLRS and research at the university. "The reviewers generally recommend stronger integration of research and development at the HLRS into the scientific setting of the university. Consequently, they particularly welcome the intention of the University of Stuttgart to establish a center for simulation technology that can collaborate interactively with the HLRS, with several professorships and an internal funding scheme."[247]

Both recommendations shifted the problem to the university, appealed to its sense of pride, and reassured the university administration that they were actually doing the right thing. The overall review of the HLRS was positive; the requirements of the Science Council had been successfully implemented; the work of the steering committee was also positively reviewed; and the users were satisfied. In addition, the establishment of the HLRS had advanced the development of scientific computing in Germany.

All that remained was to create a slightly simpler organizational structure and to ensure (this bit was for the benefit of BMBF readers) that the state would guarantee the stability of this German high-performance computing center. “The tasks of the HLRS cannot be taken up by a commercial organization but must be carried out by the public sector. As the mandate applies to all of Germany, the main responsibility should lie with the state and not with a single university.”[248]

Some readers may well have caught their breath at that final sentence of the report. No sooner had they grasped that outsourcing was being discouraged than the report took the next crucial step. A “mandate that applies to all of Germany” would logically have implied transferring the main responsibility for the HLRS to the federal government. However, the report ended with the proposition that the HLRS be more closely linked to the scientific research at the university and to the budgetary resources of Baden-Württemberg. Tossing in surprising recommendations is part of the art of reviewing.

Grid computing extends a lifeline

On sober reflection, high-performance computing in Stuttgart was in an awkward position. The further development of the HLRS appeared to be assured by any of three options: The uneven road to a network of German supercomputing centers had not been bypassed by global metacomputing. Experiments had confirmed that. In addition, the reviewers had strongly advised against organizational or budgetary outsourcing. Finally, the recommendation to link university research more closely with the HLRS was unrealistic in view of the very modest demand for courses in scientific computing, and thus was vetoed by users.[249]

At the same time, despite all the planning, high-performance computing in Stuttgart had been given back to the state to organize. Which would mean dealing with Karlsruhe.

This had already been the case in the early 1980s, during the lengthy and cumbersome process of procuring the Cray-1. At the time, the issues were how to keep the lines of demarcation clear and how to supplement federal funds for a supercomputer in such a way that each university could purchase a vector computer. Two decades later, it had again become necessary to cooperate with the competition – this time, however, not for purposes of drawing lines but with the incomparably more demanding goal of interconnecting.

It was certainly worth a try, especially considering that conditions could hardly be worse than before. The Science Council strongly supported the effort and was keen to coordinate.[250] The BMBF was working to promote a nationwide high-speed network. And UNICORE now provided a uniform interface for computing resources that would enable seamless computing even in Europe. It should also be possible to network university computers, one of which would be located in Karlsruhe, and the other in Stuttgart.

Baden-Württemberg approached the Science Council with a new proposal for establishing a high-performance computing facility in Karlsruhe and Stuttgart – and was promptly dismissed. In a statement on the proposal issued in May 2002, the council made a few remarks about its own coordinating role and competence before delving meticulously into all aspects of the proposal: planned areas of application, intended computers, participating institutions, potential users, project management, location, costs, and time planning. Everything was listed in detail. Following over twenty pages of conscientiously regurgitated proposal-ese, the conclusion came as a shock: the council doubted whether a linked computer system would enable Stuttgart and Karlsruhe to achieve the stated goals. "With computing capacity distributed between two locations, processing particularly demanding applications … is only possible to a limited extent." Despite the interesting and desirable prospects for "heterogeneous grid computing," a network of high-performance computers was still far from being a means of providing the highest computing capacity. "It remains essential for users to be able to take advantage of preprocessing, processing, and postprocessing at one location." The Science Council was even less inclined to link two different computer architectures in Baden-Württemberg. This was something that could perhaps be realized at the national level, but not in one state. In any case, at the moment maximum computing power could only be provided according to the "one-location principle." Certainly, the connection between the two subsystems in Karlsruhe and Stuttgart was an adequate technical prerequisite for a distributed file and batch system, but "high latency in networks remains unavoidable."[251]

The qualified success of the metacomputing experiments between Stuttgart, Pittsburgh, and Albuquerque at the end of the 1990s still limited the scope for action, even for the short connection between the two top computers in Baden-Württemberg. The reference to the high latency times was unmistakable for insiders. But the Science Council cleverly combined it with a recent and powerful semantic shift: "metacomputing" now exclusively denoted connecting computers specifically to expand total computing capacity.[252] The much broader term "grid computing," on the other hand, was used to describe all forms of distributed computing; it did not necessarily have to be high-performance computing. For the Science Council, too, generally speaking a computer network offered the possibility of differentiating computing centers according to function, in other words, dividing the labor.

The rejection of the proposal nullified much of the preparatory work, but it also had advantages. First, it made clear to the

zurück

Baden-Württemberg ministry in charge that the Southern German high-performance computing environment could not be shaped by forcing connections. Second, it relieved the universities concerned of the burden of an often unprofitable collaboration. And third, the Science Council again shifted the high-performance computing problem to the federal level in a way that might at least be attractive for Stuttgart. The HLRS had already been a national high-performance computing center for several years. The accelerated further development of the D-GRID project, which was intended to secure access to the computing capacity needed somewhere in Germany for the whole of Germany, made it possible for a center like the HLRS to specialize. That might mean redoubling its efforts in CFD, which was of interest to industry. Or perhaps it might expand virtual reality, which had already been tested with Sandia,[253] into a specialty area with a focus on innovative grid development for engineering applications.[254]

The questions were not easy to answer and would require additional conceptual effort. There was nothing wrong with the German science policymakers' enthusiastic embrace of grid computing. But the Science Council's rejection of the Baden-Württemberg proposal made it clear that the situation could careen out of control. If any user was able to access any computer anywhere in Germany, centers would lose their visibility, even high-performance computing centers. The high-performance computing community ran the risk of being marginalized by the top of the pyramid to which it had long aspired, and of making nothing more than an indifferent contribution to Germany's computing achievements.

A few basic points now needed to be documented. This could be done, for example, by means of *inSiDE*, a journal founded in 2003. Friedel Hoßfeld temporized in the first issue but compensated by being decidedly outspoken in the second. Forget feasibility studies, committee work, and external reviews; instead, focus on making the most of available concepts in the field of supercomputing. Which essentially meant taking all the criticism the Science Council had lobbed at supercomputing and turning it to the advantage of the high-performance computing community.[255]

Hoßfeld reminded readers of what the council had been saying on this matter since 1995. He reiterated that the pyramid needed a top, with several national centers providing supercomputing capacity.[256] Access would have to be secured via efficient data networks, and the centers within the network would be responsible for developing and disseminating methods, instruments, applications, and training.[257] This might require a Germany-wide "load-balancing model." However, what was needed most of all was an uninterrupted spiral of innovation, with adequate phase shifts, by which the three designated supercomputing centers in Stuttgart, Jülich, and Munich could provide the German science and technology community with unparalleled computing capacity.[258]

In budgetary terms, the concept of the innovation spiral was primarily an investment spiral. It was intended to ensure international competitiveness and to keep the middle and lower performance segments of the pyramid at a safe distance from the honeypots of science policymakers, while at the same time delivering on the promise of fast, location-independent access to data and IT resources available worldwide. There was no objection to the massive investment in nationwide grid computing, as announced under the leadership of the Helmholtz Association with support from the federal ministry.[259] Nonetheless, at the same time, it was necessary to ensure that at least the top of the German computer pyramid was adequately maintained. This way, there would be no need to directly couple top computers; the federal computer centers would be able to specialize independently and expand their competencies to provide complementary services.

zurück

Users at work (2006–2016)

When Rector Dieter Fritsch presented his last report on the state of the university in 2006, he and the HLRS could look back on an eventful time. As soon as he took office, Fritsch had been commissioned to “organize 45 million euros to procure a new generation of high-performance computers.” At the same time, he heard rumors that “the Stuttgart location could no longer be maintained, and that the sizable competition in the Baden area had already set course to build the new computer in Karlsruhe.” For Fritsch, it seemed the battle was “actually already lost.”[260]

Retiring rectors are afforded the courtesy of publicly recalling the highlights of their tenure. If it serves the cause and does not risk damaging their own reputation afterwards, they may even reveal various and sundry details from the kitchen cabinet of university politics. Thus did a visibly relieved Fritsch describe his inaugural meeting with then minister-president Erwin Teufel. To Fritsch’s surprise, Teufel made clear that the state might still be willing to finance a new supercomputer in Stuttgart “if a suitable, cross-university computer network concept were forthcoming.” However, Fritsch was also informed that the difference between a normal and a cutting-edge high-performance computing center amounted to 30 million euros. Moreover, this difference could, as it turned out, only be eliminated with the help of a high-performance computing competence center in Baden-Württemberg and with “considerable political skill and tactics.” In the end, the University of Stuttgart did receive a top high-performance computer, and Karlsruhe had to make do with standard equipment. Stuttgart’s new computer could even be visited in

zurück

a building constructed especially for the HLRS, Fritsch noted with satisfaction. “The University of Stuttgart has not only retained an excellent infrastructure for research and development, but we have also been able to further enhance our excellent reputation worldwide in the course of expanding the three German high-performance computing centers!”[261]

Taking it easy

The new buildings at Nobelstrasse 19 on the Vaihingen campus gave computing at the limit of computability in Stuttgart a stable university address.[262] Now the HLRS, which had bounced between commercial enterprises and the subdivision of the computing center since its founding, was organizationally embedded in the university as a “key facility”[263] (Fig. 12). However, computing centers are not permanently cast in concrete and then left to their fate. In planning, the architects in charge of the university building authority had to take into account the far-reaching changes that had occurred in the architecture of high-performance computers in the 2000s. The new machines that would reside in the HLRS had different space requirements and, in particular, an appetite for energy that far exceeded that of their predecessors.

Previously, it was safe to assume that, with increasing performance, space and energy requirements remained more or less constant. But at the turn of the millennium, manufacturers of high-performance computers decided to use standard components with air cooling for cost reasons. This did indeed reduce the energy demands of individual components. However, since a larger number of them were used over the entire system, processing power and energy consumption increased in lockstep. This principle had been demonstrated by the projects of the American

zurück

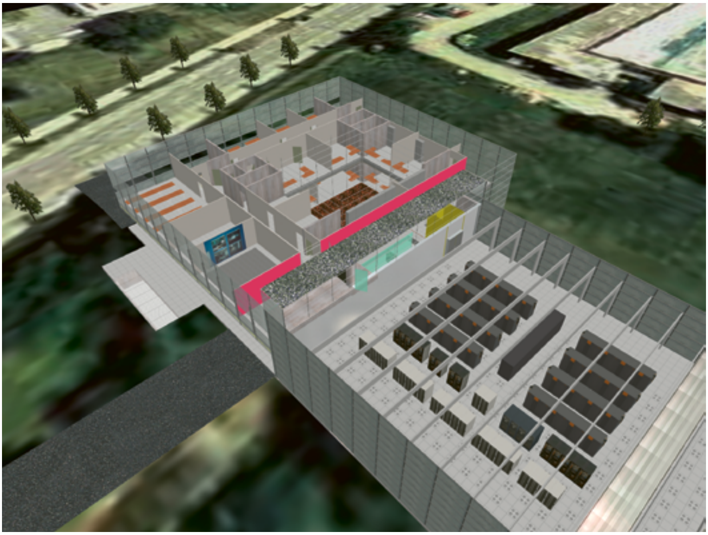

Fig. 12: Demonstrable transparency, 2003: simulating a "key facility."

Accelerated Strategic Computing Initiative as well as by the Japanese Earth Simulator project.[264]

At the new HLRS building's "topping out" ceremony, Klaus Schmiedeck, head of the university's building office, explained that the outer construction invited visitors to "take it easy." But what comes across as "easy" is typically quite challenging. Apart from the structural design, ventilation and cooling were a "major project" when the "computer house" was built. The electrical design also has to be done differently when the building's occupants "draw a lot of electricity," the *Unikurier* quoted Schmiedeck as saying.[265] But advances in computer architecture ruin even the most careful planning. The computing power of supercomputers continued to increase, each new generation of machines in turn driving up demand for electrical energy for operation and cooling.[266] For Stuttgart to remain in the rat race at the limit of computability, ever greater investments in energy technology and concrete had to be made. In 2010, a special building was commissioned from the university building office to meet the energy requirements of the anticipated supercomputers. The infrastructure building located right next to the HLRS main building was equipped with an ingenious cooling system: "A total of eight transformers ensure uninterrupted power supply and cooling. The cooling water itself is fed into the building via a separate service floor and distributed by means of a highly efficient cooling network. Conversely, the exhaust air from the computers is used to supply the building with heat."[267]

The transformers and pumps of the new infrastructure building were no trivial matter to be discussed only with the building office and dealt with only by the HLRS technical personnel. Rather, the new materiality of the supercomputers and the associated infrastructure was integrated into the aesthetic concept that determined the architecture of the new HLRS buildings from the very beginning. The guiding principle of this concept consisted in staging the compatibility of heterogeneous forms and well-placed colors.[268] Together with artist Harald F. Müller, the building office developed its own color concept for the exterior design of the infrastructure building.[269] "Especially on sunny days," states a flyer from the university building office, "a glittering play of colors ensues between the iridescent gold ventilation slats and the grass-green, smooth concrete surfaces."[270] Here, the idea was to create an impression of lightness. The same applied to the design of the latest HLRS building. Inaugurated in 2017, the new training center building echoed the "filigree construction"[271] of the research building but as a "variation of the original."[272] The surfaces of the classroom in the inner building – known as the Rühle Saal – were "monochrome white throughout," while a "blue-tinted 'glass band'" was intended to lessen the intensity

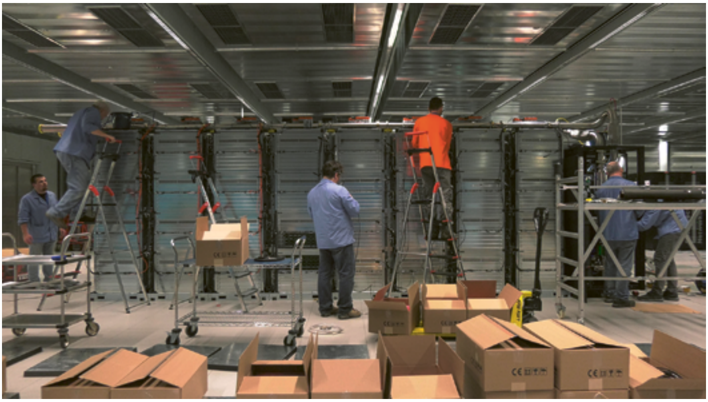

Fig. 13: Unpacking, 2020: hard at work installing the HAWK.

of the light coming from outside. The idea was to create a "quiet, open 'conceptual space,'" making the outside world "only dimly visible" and thus allowing "no distraction."[273] As a contrast to the "quiet atmosphere" inside, "randomly"[274] arranged cubes were positioned about the atrium. Here, too, the architects intended a play on references: "Randomness, which must be calculated and yet eludes calculation, which technology should protect us from and yet which arises out of technology – both converge in the connection between atrium and high-performance computer."[275]

The design principle applied all the way down to the basement of the computer rooms and even to the computer housings (Fig. 13). Their casings were decorated with graphics from HLRS simulation applications. Virtual flow patterns and automobiles covered the long black fronts of the computer clusters. What looked to be simple metal cabinets from Cray, NEC, and HPE in the windowless, air-conditioned machine rooms were actually sophisticated computer systems, which the HLRS decorated in strict alliteration with the names of endangered species – from "Her-

zurück

mit" (Cray XE-6, 2010) to "Hornet" (Cray XC 40, 2014) to "Hazel Hen" (Cray XC 40, 2015), and "Hawk" (HPE Apollo, 2020).[276]

The architecture and design of the new buildings constituted an act of self-assertion. The staging of the computers and the art and architecture reflected the desire to create an overall appearance, a uniform corporate design for the HLRS.[277] The new buildings were intended to embody and radiate the core identity of the center. They formed the symbolic capital that could be used profitably to operate high-performance computing in Stuttgart.

With Gauss to Europe

The new HLRS building, fêted by outgoing rector Dieter Fritsch in 2006, would hardly have become a long-term reliable operator of high-performance computing in Stuttgart, Baden-Württemberg, and Germany if the new location on the Vaihingen campus had not at the same time been transformed into the hub of a new science policy program. Despite his satisfaction with the rescue of the high-performance computing location, even Fritsch was not sure how all this activity was going to continue. The only thing he thought certain was that it would be quite expensive. "Rumor has it that the next generation of high-performance computers will cost 200 million euros," said Fritsch. "No federal state will be able to shoulder this [burden]," he predicted: "It would have to become a European project."[278]

Surprisingly, Fritsch betrayed little concern about this eventuality. First, it was the problem of his successor, Wolfram Ressel; and, second, the prospects for successful strategizing were much better than in 2000. Indeed, the new rectorate introduced far-reaching structural reforms in an (ultimately vain) attempt to position Stuttgart as a university of excellence.[279] Third-party

zurück

funding saw record-breaking rates of growth, an entrepreneurially oriented administration was installed, budgetary freedom was increased, and international ambition was evident.[280] The university's own self-profile described characteristics that set it apart "from other technically oriented universities in Germany" and made it "more competitive internationally."[281] As a "research university with an engineering and natural science orientation," Stuttgart now focused on cross-disciplinary topics such as simulation, which had science policy appeal, brought together many disciplines and university start-ups, and also were attractive to industry.[282] So-called Bologna-compliant teaching – an educational standard that applied across the EU – could also be "virtualized," that is, moved into digital space, with good justification.

The HLRS was a perfect building block for the structural realignment of the university. As one of the federal government's three high-performance computing centers, with close ties to industry and long-standing connections to the United States and Japan, the HLRS in Stuttgart now took on the role of an infrastructural beacon.[283] It shone all the brighter because the HLRS also became one of the new revolving doors for the university for European-funded projects and for highly visible projects financed by third parties. Unlike German priorities, the European initiatives of the 2000s used grid computing to provide the highest computing capacity.[284] Calculating at the limit of computability had become an essential component of applications-oriented European science policy. However, the focus of European initiatives was no longer on establishing a European supercomputer industry, as in the 1990s. Here, too, interest in high-performance computing had shifted from hardware to the field of application. The new funding schemes emphasized computer-based science and simulation.[285]

Shifting the national innovation spiral in high-performance computing to the European level was impossible without incurring political fallout. The cooperative model of the three German high-performance computing centers based on a network of specialized centers was too heterogeneous for negotiations at the European level. According to the ministry's instructions, German supercomputing in Europe should speak with one voice. What was required was a form of cooperation that ensured consensus among the centers and could credibly represent them to the outside world in European-level negotiations. The number of strategy papers, memoranda, workshops, and meetings increased again. Between Jülich, Munich, and Stuttgart, the difficult question of how to achieve specialization of German high-performance computing operations while keeping pace with European competition had to be resolved.

The problem was now one of several years' standing, and the concepts had been reworked again and again through updated sets of slides and reports. August 2005 saw the publication of a study, commissioned by the BMBF, on petaflop computing in Germany within the context of the European Research Area.[286] In summer 2006, the Reuter Commission worked on arguments for establishing a strategic alliance and on creating a Gauss Centre for Supercomputing (GCS), which occurred in April 2007.[287] At the same time, a memorandum of understanding for a Partnership for Advanced Computing in Europe (PRACE) was signed in Brussels. In July 2008, another memorandum of understanding was signed to establish the Gauss Alliance, followed in September by an administrative agreement on the GCS between the BMBF, Baden-Württemberg, Bavaria, and North Rhine-Westphalia.

The new European politics of consensus reduced the competition between the three federal high-performance computing centers. The fact that the pyramid that had been invoked for decades had become a tool for taming domestic German competition is evident from widely circulated statements contained in the strat-

zurück

egy papers of the German Science Council. The papers were also approvingly quoted in the Gauss Centre's PowerPoint presentations and strongly echoed its own arguments. For example, according to a statement of the GCS chairman in a 2011 strategy paper, what was needed was more high-performance computing competence networks and a waiver of demand-fueled control of high-performance computing use through fees. Essentially everything had to be coordinated – cooperation with Europe within the framework of PRACE, procurement planning, infrastructure, user access and support, and even development activities.[288]

Safeguarding competition, cooperation, alliances, and mutual assistance clauses drove the tier 0 level of the pyramid in the second half of the 2000s toward an era of petaflops. The German high-performance computing business flourished in the new configuration and, in fact, ran like clockwork, whereas "computing power, as usual, [was] still best in the USA."[289]

Virtual users, and users in virtual reality

Being embedded in the university and participating in European initiatives had once again created stable conditions for the local operation of high-performance computing in Stuttgart. Regional and European research programs brought a constant stream of new projects and new users to the HLRS.[290] Traditional user groups, such as the automotive industry, were easily integrated into the new structures.[291] The hww GmbH, founded in 1995, once again emerged a model case worthy of imitation for cooperation between data centers, science, and industry.[292] Even the process of procuring world-top-class computers lost its terror.[293] Manufacturers came to realize that for them, too, the emphasis was on performance and cooperation, not old-boy networks.[294]

zurück

Under these conditions, Stuttgart's specialization in high-performance computing could be marketed as a unique feature of the location.[295] The HLRS annual report, which was designed more like an exhibition catalog than the activity report of a computing center, began to bring users and their projects into the limelight. The projects were rendered in generally understandable digest form, and awards and publications were listed. Glossy photographs showed users as attentive listeners in introductory courses on parallel programming, CFD, and visualized simulation.[296] Alternatively, they were shown seated in cars wearing 3D glasses and driving through simulated streets. Stuttgart's expertise was particularly well represented by users who immersed themselves in the virtual aspects of their simulation in the CAVE.[297] The CAVE was a walk-in room made of acrylic and glass for 3D projections and virtual reality simulations. It impressively symbolized the continuity of Stuttgart's efforts in cooperative working environments for visualizing simulations since the 1980s.[298] In the CAVE, users could move through the virtual reality of their simulations and manipulate different variants of their models (Fig. 14). Observed from the outside, you might be forgiven for thinking that it wasn't necessary to sit at the HLRS computer to experiment with the models.[299] In images of the CAVE, the supercomputers are just as invisible as the work done by the staff at the HLRS, who programmed the algorithms and software for the CAVE's applications.[300]

The new plan for including heterogeneous user communities at the HLRS required an operating model that could differentiate users and provide them with the resources they needed.[301] Developing a Porsche auto prototype or calculating the energy output of a power plant ultimately demanded more and differently structured capacity than the test run of a programmed algorithm or simulation of a hip operation. The HLRS "operating

Fig. 14: The view from within: users in the CAVE of the HLRS.

system," which was used to set up and manage access to computer resources and services, had to be reconfigured. But doing that would require breaking old habits of thinking about how to operate high-performance computing. If the aim was to cover the entire gamut of user projects and also calculate smaller jobs, then the question arose whether increasing overall utilization of peak capacity was still a sensible target. Or whether it still made economic sense to reserve a node or processor core exclusively for a single application.[302] Was it really necessary to spring the full complexity of high-performance computing on occasional users who may once have attended a training session at the HLRS? Or would it perhaps suffice to have a web platform that could take over job management of non-computation-intensive applications? If required, such a system could even enable virtual access

zurück

to the computers and visualization programs at the HLRS from anywhere in the world.[303]

The basis of the new HLRS operating model was individually tailored packages of computing capacity, methods, software, and consulting services spelled out in service-level agreements.[304] A service-level agreement covered the entire course of a user project at the HLRS, from development to execution and completion. The agreement included the required hardware and software resources, the availability of services, business processing, data management requirements, and contractual terms and conditions.[305] With the introduction of service-level agreements, the use of high-performance computing resources became mainstream. No longer did one have to justify the need for supercomputers and simulation environments in detail on the basis of specialized scientific requirements. The HLRS had established a new business model in which high-performance computing had become a generally available commodity whose specifications could be individually negotiated and flexibly adapted at any time.

Limits to growth

In an article published in 2016, HLRS director Michael Resch reflected on current perspectives in high-performance computing.[306] He saw the field facing fundamental changes and attempted to sketch the developments of the coming years. The prognosis was clear: as stable as the specialization in high-performance computing achieved by the university, the GCS, and the connections to Stuttgart's industries might be, hardware performance was approaching its limits. If you wanted to be pessimistic, you even had to anticipate the end of Moore's law – the observation (axiomatic since the 1960s) that computing power will double every two years.

This slowdown had not happened overnight, but it had become evident only over the course of the last few years. In his article, Resch was forced to add a little history to his forward-looking perspective. The fact that processor clock frequencies would not increase indefinitely over time had long been anticipated.[307] Parallel computer architectures had made it possible to ease the impending bottleneck for many years. However, an analysis of the Top500 ranking showed that the slowest computer systems on the list had been left behind since 2009/10 and could no longer keep pace. The latest figures even showed that now the entire field would have to consider a flatter growth curve.[308]

How should global computing deal with the coming limits to growth? According to Resch, there were two different paths to follow to make any significant progress: the “thin core” concept, and the “fat core” concept.[309] The thin core concept focused on hardware developers placing as many simple processors as possible on a chip in the tradition of parallel computing. This solution did not require high clock frequencies, and mass production was not even particularly demanding. In principle, the thin core concept would make it possible to increase clock frequencies even further and to integrate even more processor cores to build a faster supercomputer. However, as Resch observed, manufacturers were actually pursuing a different path. They were applying the thin core concept to standard simulations, and thus were abandoning the high-performance computing market.[310]

The fat core concept followed a classical approach to high-performance computing architecture. Speed was achieved through increased architectural complexity. The individual processor cores were given additional functions, which transformed each “core” into a “highly tuned architecture with a number of sophisticated features.” These highly tuned processor cores consequently had to be programmed in an appropriately sophisticated

zurück

way. Manufacturers who used vector computer architectures for this purpose had to reckon with especially high prices. The fat core concept thus meant additional costs for the operators of high-performance computing centers, both for the processors and the programs.[311]

No matter how you looked at it, it was clear that it was becoming more difficult to continue to grow computing power at the same rate.[312] The slowdown in the growth of hardware features was no temporary phenomenon but a reversal of trend. For operators of high-performance computing centers, this development once again raised fundamental questions. Under such circumstances, three strategic options presented themselves: First, energy costs had risen so sharply in previous years that an improved energy balance might be achieved through use of alternative cooling systems and by less expensive generation and more efficient recycling of waste heat. This would result in significantly lower operating costs and would indirectly benefit the overall scope for investment. Second, high-performance computing directors would have to use their existing systems to develop next-generation systems and, together with suppliers, use previous models to simulate future capacity. Third, high-performance computing system users had to become more directly acquainted with the architectural requirements of their programs. Users therefore had to be brought back to the machine and involved in the programming.

These three strategic options had the advantage that they were not mutually exclusive. However, they also had the disadvantage that a high-performance computing center desiring to continue to compute at the limit of computability had little choice but to develop basic concepts simultaneously on three different fronts, and therefore to seek yet another new configuration.

A history of reconfiguration

The history of supercomputing in Stuttgart has yet to reach its end and will continue to produce surprises in the future. Our study in the history of technology has shown how the development of supercomputing was punctuated by interruptions, crises, and new starts in local operations. The interaction between university and industry, between federal and state governments, and between machines and users gave rise to a game whose rules had to be constantly adapted. Again and again, new connections were made and further boundaries delimited. Not infrequently, what had been painstakingly configured became obsolete again after only a few years and had to be reconfigured for technical, operational, scientific, or political reasons.

When, at the end of the 1970s, it was assumed that computing centers were a thing of the past, a reconfiguration based on the somewhat outmoded first vector computer from Seymour Cray was successful. After lengthy negotiations, it was purchased to replace the old mainframe machine, though, in fact, it never really could.

The most important elements of a new Stuttgart supercomputing culture became apparent in the mid-1980s. This time, a brand-new machine was installed. It was completely unclear how it was to be operated and who would be able to exploit its capacity. Both issues had to be resolved in the following years under the watchful eye of a critical public.

In the long run, there was no stopping computing at the limit of computability in Stuttgart. Consequently, toward the end of the 1990s, reliance grew on a heterogeneous machine fleet and a diverse supercomputing network to secure the exclusivity of the hub's offerings. In conjunction with the national supercomputing centers, the HLRS gained a number of strategic advantages. It was one of the computing centers on whose expertise federal science policymakers relied. Stuttgart was able to expand its specialization in the field of CFD, attract European projects, and consequently become an advertisement for its own university.

Throughout the course of its history, the HLRS has become increasingly visible. Interactive simulation techniques, in-depth user training, transparent architecture, professional marketing, and standardized user contracts also helped to operationally stabilize the center in the 2000s. However, this success did not make supercomputing a foregone conclusion. Rising energy costs, the current growth slowdown in processor performance, and new computer architectures ensure that work will continue in Stuttgart on new configurations for high-end computing at the limit of computability.

zurück

Acknowledgments

Even short history of technology studies like this one depend on support from many corners. The result at times is productive confusion among the authors and their respondents. The discussions in the History of Technology Research Colloquium and Seminar at ETH proved highly stimulating. We are also grateful for the patient tutoring in high-performance computing that we received in Stuttgart. In the course of our exchanges, we were informed, for example, that computing in Stuttgart was a big deal long before the HLRS. We were also told that it wasn't a simple matter of "history": you really had to understand a bit of the technology; the rest was just ideology. One person even said that high-performance computing was already obsolete in the last century. We have striven to put these opinions to good use as well and to arrange them historically. At times when we thought we had finally understood something, new surprises awaited us. Without the generous, constructively critical support of HLRS management, we could never have completed our project. We were supplied with extensive source material that people were delighted not to have to go through themselves. We were also permitted access to untapped material in the University of Stuttgart Archives, and we benefited from many suggestions. We would like to extend our heartfelt thanks to the following people (listed, atypically, in order of their first names) for their support: Agnes Lampke, Alexandra Hees, Andreas Kaminski, Christian Ritter, Daniela Zetti, Erich Projer, Erika Fischer, Harald Atmanspacher, Henrike Hoffmann, Joachim Buhmann, Jutta Sauer, Michael Resch, Michele de Lorenzi, Norbert Becker, Rachele Delucchi, Rainer Klank, Simone Roggen-

zurück

buck, Uwe Küster, Wilfred van Gunsteren. We are highly indebted to Giselle Weiss and her careful reading of the German edition, her wonderful translator skills, and her elegant prose.

January 2021 *D.G.* and *R.W.*

Notes

1 Universität Stuttgart 2003, p. 9. For a critical discussion, see Hoffmann 2004. On Argyris, see Brand 2012; Doltsinis 2004; Ingolf Grieger, John Argyris, in: Stuttgarter Unikurier, No. 93, April 2004.
2 Effenberger 2020.
3 Roland Rühle called the configuration he helped design the "Stuttgart model." Roland Rühle, Das Stuttgarter Modell. Gemeinsame Nutzung eines Höchstleistungsrechners durch Universität und Industrie, Informationsveranstaltung der Universität Linz zum Thema "Supercomputing für Oberösterreichs Industrie und Wissenschaft," 15. 6. 1989 (Slide set) (Küster, private collection). On the concept of configuration, see Gugerli and Tornay 2018.
4 What Z1 to Z21 were about, and why Z22 was considered the first machine is beyond the scope of this discussion.
5 Universität Stuttgart, Jahresbericht des Rechenzentrums 1971, p. 6 (UASt).
6 Universität Stuttgart, Jahresbericht des Rechenzentrums 1972, unpaged (UASt).
7 Universität Stuttgart, Rechenschaftsbericht des Rektors 1975/76, p. 246 (UASt); Universität Stuttgart, Rechenschaftsbericht des Rektors 1976/77, p. 242 (UASt).
8 Universität Stuttgart, Jahresbericht des Rechenzentrums 1976 (Foreword) (UASt).
9 Harms and Meuer 1995, p. 104. See also Ergebnisprotokoll des Gesprächs mit Prof. Haupt in Aachen am 4. 10. 1977 (UASt); Universität Stuttgart, Rechenschaftsbericht des Rektors 1977/78, p. 303.
10 Deutsche Forschungsgemeinschaft, 12. 5. 1978, Stellungnahme der Kommmission für Rechenanlagen zur Anfrage des Landes Baden-Württemberg vom 26. 7. 1977 bezüglich Erweiterung des Regionalen Rechenzentrums Stuttgart im Programm zur Errichtung Regionaler Rechenzentren (Az.: 375224/2/77), 12. 5. 1978 (UASt).
11 Rechenzentrum der Universität Stuttgart, Zukunftsrichtungen der Hardwareeplanung an der Universität Stuttgart. Integrierendes Konzept zur künftigen Versorgung der Universität Stuttgart mit verteilter Rechenkapazität vom 15. 3. 1979 (UASt): "In the last two years, the University of Stuttgart has submitted various proposals to the state of Baden-Württemberg, including:

zurück

1. Proposal of 21. 7. 1977 for the general expansion of the regional computing center, including procurement of a CYBER 175 to complete the 2nd expansion stage as provided for in the expansion plan of 1. 9. 76 (following many rounds of negotiation, the federal government declined to approve this proposal for financial reasons). 2. Proposal of 22. 1. 79 to expand the CYBER 174 system, as a reduced alternative to the unapproved proposal 1. S[ee] above. 3. Proposal by the Institutes for Nuclear Technology and Energy Systems (IKE) and Statics and Dynamics (ISD) for one computer each the size of a VAX 780 or Prime 750. 4. Proposal by the computing center (originally dated 21. 12. 76) and the engineering sciences group dated 7. 3. 77 for a mini-computer to be used primarily as a concentrator for interactive operation. For scheduling reasons, the computing center included its part of the proposal in the application for regional follow-up funding in November 1978."

12 Reinsch 1982, p. 171. According to the 1973 annual report, the machine operated in three shifts, 359 days a year. Nevertheless, there were waiting times of up to two days for jobs requiring a few minutes' computing time. See also Universität Stuttgart, Jahresbericht des Rechenzentrums 1973 (UASt). Similarly, Universität Stuttgart, Rechenschaftsbericht des Rektors 1972/73, p. 102 (UASt): "The computing center facilities are nowhere near up to the capacity requirements of users despite triple-shift operation and weekend work." The strength of the CDC mainframe computer, located at RUS, lay not in computing at the limit of computability, but rather its factory-like organization of remote-access computer processing. The machine integrated as many different types of applications as possible, not the most spectacular applications. In the 1970s, reference to triple-shift operation was a euphemistic way of saying that a company could not grow to meet demand.

13 Reinsch 1982, p. 172. Even earlier, Universität Stuttgart, Rechenschaftsbericht des Rektors 1978–1980, p. 170 (UASt): "As a result of the critical situation, our users – helped by various sources of funding – have created minicomputer capacities." Sometimes these "mini" computers were quite impressive (PDP, Prime, and VAX).

14 Reinsch 1980. Nearly a decade later, the third comprehensive IT plan described this reorganization as a process of "diversification and decentralization." Ministerium für Wissenschaft und Kunst Baden-Württemberg 1989a, p. 29 f. The first comprehensive IT plan encompassed 1971–1973 and forecast up to 1975. Work on the second comprehensive IT plan did not begin until 1977. It encompassed 1979–1985 and was issued in 1980. Ministerium für Wissenschaft und Kunst Baden-Württemberg 1980, p. 2. It is thus obvious that between 1973 and 1977 only minor changes in planning were needed,

but that the need for planning increased toward the end of the 1970s. The difficulties in obtaining a replacement for the Stuttgart CD 6600 point in the same direction.

15 It was well known in Stuttgart that Japanese regional computing centers were striving to maintain the advantageous economies of scale of large facilities. Reinsch 1980, p. 79.

16 Das Rechenzentrum, 2, 1978, pp. 59 and 60 (Geleitwort des Herausgebers).

17 Reinsch to the Rektoramt der Universität Stuttgart, Dezernat Finanzen, 3. 5. 1978 (UASt): "We herewith send you the 1978/79 capital expenditure proposals and request that you forward them to the DFG for funding from the regional program." In fact, the University Construction Act (HBFG) would cover a large part of the further cost of expansion. The federal government had been supporting state university expansion projects under this law since 1969. Between 1973 and 1978, for example, a total of 73 data processing systems were purchased at Baden-Württemberg's universities under the HBFG at a modest cost of 22.2 million marks. See also Ministerium für Wissenschaft und Kunst Baden-Württemberg 1980, p. 8.

18 Eberhard Staiger, IBM Deutschland, to Karl-Gottfried Reinsch, 11. 4. 1978 (Vermerk: Benchmark IBM System 3033) (UASt); U. Löffler and W. Biela, Control Data GmbH, to Karl-Gottfried Reinsch, 22. 2. 1978 (UASt); Hertweck et al. 1979.

19 Universität Stuttgart, Rechenzentrum der Universität Stuttgart, Zukunftsrichtungen der Hardwareplanung an der Universität Stuttgart. Integrierendes Konzept zur künftigen Versorgung der Universität Stuttgart mit verteilter Rechenkapazität, 15. 3. 1979 (UASt).

20 Reinsch 1982, p. 101.

21 Karl-Gottfried Reinsch, Was heisst und zu welchem Ende betreibt man einen Grösstrechner. Gewisse notwendige Elemente eines Konzepts. Stellungnahme zu Aspekten der Informationsverarbeitung in Baden-Württemberg (Dok. Nr.: 79/50, RD 16/500-1), October 1979, Stuttgart, p. 7 (UASt). The reference is to an address given by Friedrich Schiller at Jena University on 26–27 May 1789: "What is and to what end does one study universal history?" [trans.].

22 Reinsch recognized this solution while studying large regional computing centers in Japan. Here, he observed, was where new network services and architectures were being created which formed the "basis for a computer network." Computers could thus be used as a means of "function sharing," making it possible to alleviate "pre-programmed bottlenecks for computing centers using a wide range of user software" and to offer "support and program maintenance." Reinsch 1980, p. 77.

zurück

23 The threat posed to RUS by the decentralized institute computers did not imply a decline in demand for central computing capacity. In the late 1970s, a number of disciplines not necessarily associated with computers also began to make use of mainframe capabilities, for example, corpus linguistics, sociology, and cliometrics. Although these disciplines were not represented at Stuttgart, the disciplinary structure of the user community nevertheless became more diverse, as attested by the varied program library at RUS. D. Kirchgraber (Rechenzentrum der Universität Stuttgart), Die Programmbibliothek. Allgemeine Übersicht, 1978 (UASt). The establishment of a computer user committee chaired by Vice-Rector Franz Effenberger during the rectorship of Hartmut Zwicker also suggests a heterogeneous user landscape. Effenberger 2020, p. 73.

24 Universität Stuttgart, Rechenschaftsbericht des Rektors 1978–1980, p. 171 (UASt).

25 This is a topos that appears repeatedly in connection with the future design of the Stuttgart computing center – the first mention was on the occasion of the founding of the university's Regional Computing Center in 1972.

26 Universität Stuttgart, Wissenschaftlicher Grösstrechner für das Land Baden-Württemberg. Antrag zur Realisierung am Regionalen Rechenzentrum der Universität Stuttgart, Stuttgart 15. 3. 1980 (UASt).

27 Mönkediek 2009, p. 24.

28 Leimbach 2011, p. 166.

29 Sandner 2008, p. 193. This working group gave rise in 1981 to a further working group for heads of university computing centers in Baden-Württemberg, at whose initial meeting at Lake Constance the members agreed to meet regularly to discuss relevant issues ("Konstanzer Seefrieden").

30 Sandner 2008, p. 193.

31 Ibid.

32 When launched in 1976, the Cray-1's unprecedented computing power made computer history. The first computing center in Germany to operate a Cray-1 was the Max Planck Institute for Plasma Physics in Garching (1979). Hoßfeld 1984, p. 280. On the federal government's Data Processing Program funding scheme, see Leimbach 2011, p. 166. The second Data Processing Program, which ran from 1971 to 1975, was intended to provide "easily accessible computing capacity" to all areas of research and teaching at the university. In addition, the program combined the "German research foundation's computer equipment procurement program, computer procurement under the federal government's program for funding university expansion, and the program for regional large-scale computing centers" into a single program co-

zurück

ordinated by the DFG. See also: Forschungsbericht (IV) der Bundesregierung 1972, p. 98. With the end of the third Data Processing Program in 1979, procurement of mainframe computers, which were co-financed by the federal government, stagnated. See Universität Stuttgart, Rechenschaftsbericht des Rektors 1978–1980, p. 170 (UASt).

33 Effenberger 2020, p. 75. See also Natalja Krieger (2014): Station Q: Mikroelektronik und Rechenzentrum bzw. Technische Informations- und Kommunikationsdienste, www.uni-stuttgart.de/universitaet/profil/historie/campus/stationen/vaihingen/west/info_station_q1.pdf (accessed 6. 10. 2020).

34 Späth 1985. On the reorganization of science policy guidelines in West Germany around the middle of the 1980s, see Gall 2001. On the discovery of societal shaping, see also Evers and Nowotny 1987.

35 In 1985, the German Science Council noted in this regard: "Competition presupposes a certain degree of freedom of action for those involved in the competition. To be competitive, they must have the right and the opportunity to decide whether to provide individual services and to respond to the demand for resources." Wissenschaftsrat 1985, p. 7.

36 Jürgen Blum, Vorwort, in: Jürgen Blum (ed.), Höchstleistungsrechner. Anwendung. Finanzierung. Organisation. Ein Bericht über ein Expertenseminar, Stuttgart 1985 (Küster, private collection).

37 Roland Rühle, Anwendung für Vektorrechner in den Universitäten am Beispiel der Universität Stuttgart, in: Jürgen Blum (ed.), Höchstleistungsrechner. Anwendung. Finanzierung. Organisation. Ein Bericht über ein Expertenseminar, Stuttgart 1985, pp. 1–10 (Küster, private collection).

38 Rolf Theenhaus, Anwendungen für Höchstleistungsrechner bei Grossforschungseinrichtungen am Beispiel der KFA Jülich, in: Jürgen Blum (ed.), Höchstleistungsrechner. Anwendung. Finanzierung. Organisation. Ein Bericht über ein Expertenseminar, Stuttgart 1985, p. 12 (Küster, private collection).

39 Peter E. Schuhe, Organisation und Finanzierung des Konrad-Zuse-Zentrums für Informationstechnik Berlin (ZIB), in: Jürgen Blum (ed.), Höchstleistungsrechner. Anwendung. Finanzierung. Organisation. Ein Bericht über ein Expertenseminar, Stuttgart 1985, pp. 1–7 and appendix (Küster, private collection); Wolfgang Männel, Kostenverteilung für Investitionen und Betrieb von Höchstleistungsrechnern im Rahmen einer Kooperation von Industrie, Grossforschungseinrichtungen und Universitäten, in: Jürgen Blum (ed.), Höchstleistungsrechner. Anwendung. Finanzierung. Organisation. Ein Bericht über ein Expertenseminar, Stuttgart 1985, pp. 1–17 (Küster, private collection).

40 Jürgen Blum, Eine Vertriebsgesellschaft für das Höchstleistungsrechenzentrum an der Universität Stuttgart?, in: Jürgen Blum (ed.), Höchstleistungs-

rechner. Anwendung. Finanzierung. Organisation. Ein Bericht über ein Expertenseminar, Stuttgart 1985, p. 3 (Küster, private collection). The rector's office and the state Ministry of Science and Art were occupied with the organizational considerations for many months. On 22. 11. 1985, the rector's office submitted a report on the subject from the Institute for Business Analysis. On 15. 10. 1986, Blum sent a "draft of organizational considerations for operating the Cray-2 supercomputer at the University of Stuttgart and establishing a distribution company" by messenger to senior government councillor Peters at the Ministry of Science and Art (UASt).

41 The principle of freeing up resources on supercomputers was applied conspicuously as early as the mid-1960s at NASA's Mission Control in Houston. Houston, too, had input computers, backup computers, systems for graphic analysis, and some 2,000 employees, which collectively served to reduce computer processing requirements. See also Gugerli 2018, pp. 88–105.

42 This brochure is a carefully designed, unpaged booklet describing RUS. It was probably disseminated in the first half of 1986 on the occasion of the upcoming delivery of a Cray-2 in September 1986 and was kept by Jutta Sauer. The cover and publishing information are missing. Rechenzentrum Universität Stuttgart 1986.

43 A PDP 11/40 was available as an interactive graphics computer at RUS as early as 1975. Rühle 1990, p. 5.

44 On the Cray-1's front-end computer, whose procurement had sparked intense discussion among the computer user committee and was ordered not from Siemens, but IBM, see Effenberger 2020, p. 74 f.

45 A brochure published by RUS in August 1987 emphasized the variety of network access: "The computing center's facilities can be accessed via local and national networks (Ethernet, HYPERchannel, DATEX-P/DFN, EARN) using different technologies (IBM, DECnet, TCP/IP, OSI protocols)." Rechenzentrum der Universität Stuttgart, August 1987 (UASt).

46 Loebich et al. 1987, p. 653.

47 Rühle 1990, p. 4. See also Rühle 1971 and Rühle 1973.

48 Loebich et al. 1987, p. 653.

49 Computer Zeitung, 4. 3. 1987.

50 The exclusivity clause read as follows: "Naturally, we are in talks with several companies in Germany about a Cray-2, but we do not expect to make another delivery within 12 months after delivery of your Cray-2. Should we receive such requests, we will of course request your approval." Cray Research Inc. to Lothar Späth, n.d. (UASt).

51 Computer Zeitung, 4. 3. 1987.

52 The Cray-1 had to be kept running longer than planned and could only be retired in July 1987. When the Cray-1 finally did shut down in the summer of 1987, usage of the Cray-2 was at a good 30 percent and rose to about 70 percent by the end of the year. See Karl-Gottfried Reinsch, Bericht Betriebszeitraum CRAY-2, November 1986 to October 1987, p. 9 (UASt). This also led to a conflict with the building authorities, as the additional power requirement due to the overlapping operation could only be covered through the provisional installation of a transformer that was needed elsewhere and would lead either to construction planning uncertainties or additional investment costs. Universitätsbauamt to Rechenzentrum der Universität Stuttgart, 13. 10. 1986 (UASt).

53 On 20. 9. 1985, the Ministry for Economic Affairs and Technology of Baden-Württemberg sent a letter with a subject line reading "Procurement of a Cray 2 supercomputer for the University of Stuttgart, herein: Joint use of the computer by companies," along with a list of interest parties (UASt). A letter from the chancellor of the University of Stuttgart to Dr. Wilfert from the technical data processing department of Daimler-Benz dated 20. 10. 1986 begins: "I have been trying unsuccessfully for some time to reach you by phone." His motive was "to offer industry joint use of the University of Stuttgart's computing center." Daimler-Benz responded with silence. The plan to set up a distribution company failed owing to objections from the state audit office. See Kanzler Blum to Herrn Beyer, IKOSS, 3. 11. 1986 (UASt): "We are writing to inform you that, unfortunately, our efforts to install a CRAY-2 supercomputer are currently on hold. As you most certainly must have gathered from the press, the state audit office has serious reservations about such companies." The negotiations with the Max Planck Society, which began in May 1985, also failed. See Zwicker to Präsident der Max-Planck-Gesellschaft, 24. 6. 1985 (UASt). The University of Stuttgart persisted in its attempts at least until the summer of 1988. See Generalverwaltung der Max-Planck-Gesellschaft to Rektoramt der Universität Stuttgart, 20. 9. 1988 (UASt): "Since the Max Planck Institutes are now able to obtain vector computing time from other computing centers at favorable conditions, the Garching computing center is no longer fully utilized. We therefore ask for your understanding while we wait and see how the demand for vector computing time at the Max Planck Institutes develops."

54 Computer Zeitung, 4. 3. 1987.

55 Reinsch 1987, p. 103.

56 MacKenzie 1991.

57 Stürme aus dem Rechner, in: hightech 5/89, p. 5 (Sauer, private collection).

zurück

58 Wolfgang Brand, Spitzenrechner für Spitzenleistungen. Die Geschichte des Einsatzes von Hochleistungsrechnern in der Crashsimulation bei Porsche, www.hi.uni-stuttgart.de/wgt/forschung/studiprojekte/Projektseminar-Porsche-/porsche-hausarbeiten-copy/Brand_pop.pdf (accessed 23. 9. 2020).

59 Breisgaumetropole schließt sich Stuttgarter Superrechner an, in: Badisches Tagblatt, 20. 9. 1988.

60 The establishment of this company was in tune with Stuttgart chancellor Blum's call at the expert seminar for new organizational models to market Cray computing capacity. But the inception proceeded differently from the plan envisioned in Stuttgart. Blum complained to the ministry about "various points that had not yet been sufficiently clarified" (p. 1) for a computing center in the Upper Rhine region. Among other things, he regretted that a right of participation in the steering committee was seen as a quid pro quo for use of the Cray-2. This was out of the question for the University of Stuttgart. According to the university statutes, it was simply not possible for a private sector organization to have influence over the university's committees. Blum (Kanzler Universität Stuttgart) to Ministerium für Wissenschaft und Kunst, 27. 11. 1986 (UASt).

61 Christian Brauner, "Mit Hilfe des rasend schnellen Computers die lästige Konkurrenz ausbremsen," in: Badische Zeitung, 3. 2. 1989, unpaged (Sauer, private collection).

62 Ibid. (italics in original).

63 Ibid.

64 Ibid.

65 Ibid.

66 Universität Stuttgart, Beschaffung eines Höchstleistungsrechners vom Typ CRAY-3 für die Universität Stuttgart, 1989, p. 5 (Küster, private collection).

67 Ibid.

68 Ibid., p. 10 f.

69 Lang 1989.

70 Ibid.

71 Universität Stuttgart, Beschaffung eines Höchstleistungsrechners vom Typ CRAY-3 für die Universität Stuttgart, 1989, p. 9 (Küster, private collection).

72 Ibid.

73 Lang 1993, p. 126.

74 Rechenzentrum Universität Stuttgart 1994, p. 12.

75 Universität Stuttgart, Rechenschaftsbericht des Rektors Prof. Dr.-Ing. Jürgen Giesecke, 1. 10. 1990 to 30. 9. 1991, p. 37 (UASt).

76 Lang/Rühle 1992, p. 122.

zurück

77 Ibid.
78 Ibid.
79 Ibid.
80 Ibid., p. 127.
81 Schlecht 1989, p. 1.
82 Universität Stuttgart, Beschaffung eines Höchstleistungsrechners vom Typ CRAY-3 für die Universität Stuttgart, 1989, p. 9 (Küster, private collection).
83 Universität Stuttgart, Forschungsbericht 1995/96, p. 348 (UASt).
84 Ibid.
85 Gentzsch 1987.
86 Meuer 1992, p. 388.
87 Cray Research Inc. to Lothar Späth, n.d. (UASt).
88 Roland Bischoff, "Cray 2 schon voll? Ein Gespräch mit den beiden Direktoren des Rechenzentrums der Universität Stuttgart, Prof. Dr. Roland Rühle und Dr.-Ing. Karl-G. Reinsch über die Auslastung und Engpässe des Supercomputers Cray 2," in: Computer Magazin, 7. 8. 1988, p. 2 (Sauer, private collection).
89 Übelmesser1989 p. 36.
90 Ibid.
91 Universität Stuttgart, Beschaffung eines Höchstleistungsrechners vom Typ CRAY-3 für die Universität Stuttgart, 1989, p. 26 (Küster, private collection).
92 Ibid.
93 Ibid.
94 Ibid., p. 1.
95 Ibid.
96 Ibid., p. 27.
97 Ibid.
98 Meuer 1990, p. V f.
99 Meuer 1991b, p. V.
100 Ibid.
101 Jenkins 1986.
102 Meuer 1991a.
103 Ahlrichs 1993, p. 179.
104 Ibid.
105 Rolf Kuhfuss/Alfred Geiger/Uwe Küster, Supercomputer der 90'er Jahre (vertraulich), 5. 12. 1989, Stuttgart, Foreword, unpaged (Küster, private collecp tion).
106 Ibid.
107 Bailey 1992.
108 Gentzsch 1987.

zurück

109 Rolf Kuhfuss/Alfred Geiger/Uwe Küster, Supercomputer der 90'er Jahre (vertraulich), 5. 12. 1989, Stuttgart, Foreword, unpaged (Küster, private collection).

110 Ibid.

111 For a comparison between *Forbes* and the Top500, see Harald Lux, "Hans-Werner Meuer ist Herr über die Hitparade der schnellsten Rechner der Welt," in: Die Zeit, 13. 6. 1997.

112 Lax 1982, Appendix I. Information on Selected Facilities, 26. 12. 1982.

113 Wissenschaftsrat 1987, Anhang 1. Ausstattung der Universitäten mit Rechnern im Rechenzentrum, in der Bibliothek, in der Hochschulverwaltung, in den Medizinischen Einrichtungen und in den Fachbereichen/Instituten (Stand: März 1987), pp. 47–65 (Stuttgart: p. 49).

114 Dongarra and Luszczek 2011, p. 2056.

115 Dongarra et al. 1979. For the in-house history of the LINPACK Benchmark, see especially Dongarra/Luszczek/Petite 2003.

116 Harald Lux, "Hans-Werner Meuer ist Herr über die Hitparade der schnellsten Rechner der Welt," in: Die Zeit, 13. 6. 1997.

117 Dongarra/Luszczek 2011, p. 2056.

118 www.top500.org/lists/top500/1993/06/ (accessed 23. 9. 2020).

119 Universität Stuttgart, Antrag der Universität Stuttgart auf die Beschaffung eines Höchstleistungsrechnersystems als Nachfolger für die Cray-2, 7. 7. 1992, p. 5 (Küster, private collection).

120 Ibid.

121 Only one production model of the Cray-3 was ever made. It was loaned to the National Center for Atmospheric Research (NCAR) in 1993. Anthes 1994. According to the German version of Wikipedia: "Development dragged on from 1988/1989 to 1993. With the security of a contract and government backing, Cray managed to unveil a model of the Cray-3 model in 1993 at a cost of $120 million and to install it in the NCAR" (https://de.wikipedia.org/wiki/Cray-3, accessed 23. 9. 2020).

122 Reuter 1993.

123 Universität Stuttgart, Antrag der Universität Stuttgart auf die Beschaffung eines Höchstleistungsrechnersystems als Nachfolger für die Cray-2, 7. 7. 1992, p. 33 (Küster, private collection).

124 Ministerium für Wissenschaft und Kunst Baden-Württemberg 1989b, p. 194.

125 "No specific plans to achieve this desirable goal have been made known to the [research] committee. Moreover, it recommends maintaining close contact with the corresponding efforts of the University of Karlsruhe for the pur-

zurück

pose of exchanging results and experience." Ministerium für Wissenschaft und Kunst Baden-Württemberg 1989b, p. 194.

126 Rechenschaftsbericht des Rektors Prof. Dr. Heide Ziegler, 1. 10. 1992 to 30. 9. 1993, p. 25 (UASt).

127 The Institute for Parallel and Distributed High Performance Computers (IPVR) was founded in March 1989 as part of the Department of Computer Science. Universität Stuttgart, Rechenschaftsbericht des Rektors Prof. Dr. Franz Effenberger, 1. 10. 1988 to 30. 9. 1989, p. 17 (UASt). Cooperation with the computing center was clearly on the agenda. See Prof. Dr. A. Reuter to Thinking Machines GmbH, 29. 1. 1991 (Küster, private collection). The rector's subsequent financial statement reported: "Construction of the new 'Institute for Parallel and Distributed High Performance Computers' is progressing well. In the meanwhile, cooperative agreements with IBM and Digital Equipment as well as a cooperative agreement with the University of Tübingen have been concluded," in: Universität Stuttgart, Rechenschaftsbericht des Rektors Prof. Dr. Franz Effenberger, 1. 10. 1989 to 30. 9. 1990, p. 21 (UASt).

128 Rechenschaftsbericht des Rektors Prof. Dr. Heide Ziegler, 1. 10. 1992 to 30. 9. 1993, p. 25 (UASt).

129 Alfred Geiger, Neuer Rechner. Massively Parallel Supercomputing: Installation einer intel Paragon XP&S-5 an der Universität Stuttgart, in: BI. Informationen für Nutzer des Rechenzentrums, 11/12, 1992, p. 1 (UASt).

130 Gentzsch 1987.

131 Hord 1982.

132 Alfred Geiger, Neuer Rechner. Massively Parallel Supercomputing: Installation einer intel Paragon XP&S-5 an der Universität Stuttgart, in: BI. Informationen für Nutzer des Rechenzentrums, 11/12, 1992, p. 1 (UASt).

133 Ibid., p. 2.

134 Ibid., p. 1.

135 Ibid.

136 Ibid., p. 2.

137 Ibid.

138 Rechenzentrum der Universität Stuttgart 1994, p. 21.

139 Ibid.

140 Alfred Geiger, Paralleles Rechnen an der Uni Stuttgart, in: BI. Informationen für Nutzer des Rechenzentrums, 1/2, 1995, p. 4 (UASt).

141 Ibid.

142 Ibid., p. 5.

143 Ibid.

144 Ibid., p. 6.

145 Ibid.
146 Ibid., p. 7.
147 Ibid.
148 Ibid., p. 4.
149 Ibid.
150 Rechenschaftsbericht des Rektors Prof. Dr. Heide Ziegler, 1. 10. 1993 to 30. 9. 1994, p. 40 (UASt).
151 Ibid.
152 Referat für Presse- und Öffentlichkeitsarbeit der Universität Stuttgart, Forschung-Entwicklung-Beratung, March 1995, p. 185 (UASt). For a detailed account on the individual projects, see Rechenzentrum Universität Stuttgart 1994, pp. 79–87.
153 Rechenzentrum Universität Stuttgart 1994, p. 79.

154 Universität Stuttgart, Antrag der Universität Stuttgart auf Beschaffung eines Höchstleistungsrechnersystems als Nachfolgerechner für die Cray-2, 7. 7. 1992, p. 1 (Küster, private collection).
155 Rühle 1995, p. 262.
156 HPSC 1992, p. 5.
157 Wissenschaftsrat 1995a, p. 9.
158 Ibid.
159 Ibid, p. 17.
160 Ministerium für Wissenschaft und Forschung Baden-Württemberg 1995, p. III (Foreword).
161 Rühle 1995, p. 262.
162 Wissenschaftsrat 1995a, p. 17.
163 Ibid., p. 23.
164 Ibid., p. 24.
165 Ibid., p. 25.
166 Rechenzentrum Universität Stuttgart 1994, p. 1.
167 Ibid.
168 Ibid.
169 Once "the cabinet made 35 million marks available from the sale of the building's fire insurance, following the rectorate's initiative vis-à-vis Minister-President Teufel, it was finally possible to restart the proposal process with the DFG." (Rechenschaftsbericht des Rektors Prof. Dr. Heide Ziegler, 1. 10. 1993 to 30. 9. 1994, p. 37 (UASt)). The concept favored by the DFG envisaged hooking a vector computer up to a parallel computer. A "suitable form of organization" (ibid.) was then discussed in various working groups involving representatives of all partners. Einladung zur Arbeitsgruppe

zur “Erarbeitung eines Nutzungs- und Betriebskonzepts” für den Höchstleistungsrechner in Stuttgart am 23. 3. 1995 im Mannheimer Hof, Ministerium für Wissenschaft und Forschung Baden-Württemberg to Roland Rühle, 9. 3. 1995 (Küster, private collection). Subsequent to these consultations, the shareholder agreement with hww GmbH was signed on 26. 7. 1995.

170 Rektorin der Universität Stuttgart Prof. Dr. Heide Ziegler: Grußwort zur HLRS Einweihung, in: F. Rainer Klank, Der Festakt, Einweihung des Bundes-Höchstleistungsrechenzentrums, unpaged, https://elib.uni-stuttgart.de/bitstream/11682/5778/1/186.pdf (accessed 23. 9. 2020).

171 Ibid.

172 Rechenschaftsbericht des Rektors Prof. Dr. Heide Ziegler, 1. 10. 1994 to 30. 9. 1995, p. 9 (UASt).

173 Rektorin der Universität Stuttgart Prof. Dr. Heide Ziegler: Grußwort zur HLRS Einweihung, in: F. Rainer Klank, Der Festakt, Einweihung des Bundes-Höchstleistungsrechenzentrums, unpaged, https://elib.uni-stuttgart.de/bitstream/11682/5778/1/186.pdf (accessed 23. 9. 2020).

174 Harms/Meuer 1995, p. 106.

175 Ibid. In this context, the rector of the University of Stuttgart also referred to Marcus Bierich, then chairman of the board of Stuttgart-based electronics firm Robert Bosch. Rektorin der Universität Stuttgart Prof. Dr. Heide Ziegler: Grußwort zur HLRS Einweihung, in: F. Rainer Klank, Der Festakt, Einweihung des Bundes-Höchstleistungsrechenzentrums, unpaged, https://elib.uni-stuttgart.de/bitstream/11682/5778/1/186.pdf (accessed 23. 9. 2020).

176 Initially, the University of Stuttgart held 25 percent of the share capital, the state of Baden-Württemberg 25 percent, debis Systemhaus GmbH 40 percent, and Porsche the remaining 10 percent. See Rechenschaftsbericht des Rektors Prof. Dr. Heide Ziegler, 1. 10. 1994 to 30. 9. 1995, p. 9 (UASt).

177 Ministerium für Wissenschaft und Forschung Baden-Württemberg 1995, p. 102.

178 Ibid.

179 Planning for the HLRS proceeded in two stages. At each stage, fixed amounts of money were to be released by the partners and new computers were to be procured. Harms and Meuer provide a concise overview of the model: “The idea is to think big: the state has firmly committed its share of 15 million marks and another 20 million marks for a second tranche. The federal government will contribute supplementary funds of 15 million marks. The Science Council has already endorsed the federal funding of 20 million marks. Industry will contribute appropriate computers and computer use worth more than 40 million marks, as well as associated know-how, to the cooperative effort – i.e., 70 million marks to start and a total of 110 million marks in the final stage for the 'supercenter' in Stuttgart." Harms/Meuer 1995, p. 106.

180 Rechenschaftsbericht des Rektors Prof. Dr. Heide Ziegler, 1. 10. 1995 to 30. 9. 1996, p. 54 (UASt).

181 Höchstleistungsrechenzentrum Universität Stuttgart (HLRS), Richtlinien für die Organisation, die Nutzung und den Betrieb, 14. 6. 1996, unpaged (UASt).

182 Ibid.

183 Ibid. With regard to industry, debis Systemhaus was to organize the allocation of capacity, while Porsche wished to use the computers exclusively for its own purposes.

184 Harms/Meuer 1995, p. 106.

185 Höchstleistungsrechenzentrum Universität Stuttgart (HLRS), Richtlinien für die Organisation, die Nutzung und den Betrieb, 14. 6. 1996, unpaged (UASt).

186 Ibid.

187 Ibid.

188 Rechenschaftsbericht des Rektors Prof. Dr. Heide Ziegler, 1. 10. 1995 to 30. 9. 1996, p. 54 (UASt).

189 F. Rainer Klank/Ulrich Lang/Andreas Rozek, Ministerpräsident Teufel am RUS, in: BI. Informationen für die Nutzer des Rechenzentrums, 3, 1996, p. 12 (UASt).

190 Ibid.

191 Ibid., p. 15.

192 Ibid., p. 18.

193 Ibid.

194 Rechenschaftsbericht des Rektors Prof. Dr. Heide Ziegler, 1. 10. 1994 to 30. 9. 1995, p. 10 (UASt).

195 Klaus von Trotha: "Deutscher Wissenschaft in der internationalen Konkurrenz einen Spitzenplatz sichern?," in: F. Rainer Klank, Der Festakt, Einweihung des Bundes-Höchstleistungsrechenzentrums, unpaged, https://elib.uni-stuttgart.de/bitstream/11682/5778/1/186.pdf (accessed 23. 9. 2020).

196 F. Rainer Klank, Der Festakt, Einweihung des Bundes-Höchstleistungsrechenzentrums, unpaged, https://elib.uni-stuttgart.de/bitstream/11682/5778/1/186.pdf (accessed 23. 9. 2020).

197 Ibid.

198 Video der Landespressekonferenz, 30. 1. 1996 (UASt).

199 In June 1996 the NEC SX-4/32 shot to place 10 in the Top500 ranking. In June 1997 the Stuttgart T3E Cray HPE also reached place 10, and the NEC SX-4/32 retreated to place 38. www.top500.org/lists/top500/1996/06/ (accessed 23. 9. 2020).

200 According to the so-called Bangemann report: "Two features are essential to the deployment of the information infrastructure needed by the information society: one is a seamless interconnection of networks and the other that the services and applications which build on them should be able to work together (interoperability)." European Commission 1994, p. 17.

201 Ministerium für Wissenschaft und Forschung Baden-Württemberg 1995.

202 On the GMD, see also Wiegand 1994. The GMD appears to have had a long-standing, tense relationship with the federal ministry. Initially, there was disagreement over whether the GMD should be theoretically or practically oriented. In the 1990s, it served as a catch-all for the institutes of the Berlin-based German Academy of Sciences. Between 2000 and 2001 the GMD was integrated into the Fraunhofer-Gesellschaft for application-oriented research.

203 Hoßfeld et al. 1997. Interestingly, Roland Rühle was not listed as HLRS director but rather as "head of RUS." That is, the feasibility analysis treated the four institutions as equals and dispensed with federal distinctions.

204 The mandate for the study was a trifle shaky. The report's title page carried an indirect order of the "Federal Ministry of Education, Science, Research and Technology," namely, "FE Project" with the identification number "01 IR 602/9." The responsibility for the content of this publication lay with the leading institution, i.e., the Central Institute for Applied Mathematics at Forschungszentrum Jülich. The feasibility study underwent several textual revisions (see Hoßfeld 1998), and was published in the proceedings of the Supercomputing '98 conference in Mannheim, which took place from 18 to 20 June 1998.

205 Hoßfeld et al. 1997, p. 5. For further evidence of this topos in the German high-performance computing context, see Rolf Theenhaus, Anwendungen für Höchstleistungsrechner bei Grossforschungseinrichtungen am Beispiel der KFA Jülich, in: Jürgen Blum (ed.), Höchstleistungsrechner. Anwendung. Finanzierung. Organisation. Ein Bericht über ein Expertenseminar, Stuttgart 1985, p. 9 (Küster, private collection); Reinsch 1988, p. 21. The topos is usueally found in prominent places of relevant texts, for example, the Rubbia Report of the European Commission (Commission of the European Communities 1992, p. 213), in Gropp et al. 1994, p. 1, or much later still in Wagner 2009, p. 45. It was readily taken up in media history and (presumably) lobbed as a provocation at the history of ideas branch of history of science and at philosophy of science, which simply ignored it. See also Pias 2011.

206 Wissenschaftsrat 1995a; Wissenschaftsrat 1995b.

207 Deutsche Forschungsgemeinschaft 1991.

zurück

208 Hoßfeld et al. 1997, p. 7 f.

209 Ibid., p. 8 f.

210 Almond had worked at RUS since 1967 and had long been deputy director under Reinsch. “Dr. James C. Almond, PhD, born 1933, studied chemical engineering at the University of Washington, Seattle, Washington, USA. At the Computing Center of the University of Stuttgart since 1967; now deputy director.” Author credits below text: James C. Almond, ECODU (European Control Data User) in: Rechenzentrum, No. 3, 1980, p. 181 f.

211 Almond 1990, p. 186.

212 Ibid.

213 Ibid. (italics in original).

214 Ibid., p. 188.

215 Ibid.

216 Weingart 2006.

217 Hoßfeld et al. 1997, p. 9.

218 The report cited Smarr/Catlett 1992 and Khokhar 1993. But work done in Stuttgart, including that of Geiger 1993, Geiger 1994, and Resch et al. 1994, was most likely also implied.

219 Hoßfeld et al. 1997, p. 10.

220 One of the goals of the Accelerated Strategic Computing Initiative was to use high-performance computing systems to analyze data from nuclear weapons tests and possibly even replace them with computational simulation of nuclear explosions. According to an October 1995 issue of *Science* magazine, “[t]he object is to store, share, and analyze the vast quantities of data collected from the 1,027 nuclear tests conducted since 1945, from 50 years of research at the labs, and from the new experiments conducted under the guise of stockpile stewardship. All this data and processing power, say planners, should enable scientists to create virtual nuclear explosions that will model the performance of stored, aging weapons.” Weisman 1995, p. 21.

221 Hoßfeld et al. 1997, p. 10. Refocusing the attention of national funding agencies occurs through the interaction of proposals and expert opinion, or through the strategic definition of programs negotiated in committees. These negotiations are shaped largely by scientific expertise. Nonetheless, nothing prevents anyone from characterizing a funding agency’s decision not to support particular interests either as an obvious mistake or regrettable oversight, an incomprehensible choice, or even deliberate obstruction.

222 Hoßfeld et al. 1997, p. 11.

223 Ibid., p. 12.

zurück

224 It had been assumed since the 1950s that big computers offered a better price–performance ratio than small ones. This old rule of thumb was now increasingly obsolete thanks to more powerful and cheaper microprocessors. See Cale et al. 1979. It became apparent that a network of many microprocessors operating in parallel could be built much more cheaply than a significantly improved vector computer ever could.

225 Hardware-wise, network technology did not advance to the same extent as did processor speed. Yet connecting massively parallel computers relied heavily on powerful communication technologies. Software-wise, the algorithms encountered problems related to the independence of computer processes, and mathematical problems related to the independence of operations. In short, it turned out to be anything but easy to program independent, parallel computing processes. Gropp et al. 1994, p. 2 f.

226 See also overviews of MPI in Gropp et al. 1994, pp. 1–10; Simon 1995; Meuer/Strohmaier 1996. On an alternative library known as PVM (parallel virtual machine), see Geist 1994.

227 Resch et al. 1997, p. 13. For a discussion on how PACX-MPI works, see Brune et al. 1999.

228 Beisel et al. 1997.

229 Gabriel et al. 1997.

230 Michael Resch, unpublished interview, 12. 11. 2019.

231 "Pittsburgh and Stuttgart Inaugurate High-Speed Transatlantic Metacomputing," Pittsburgh Supercomputing Center, press release, 24. 6. 1997.

232 Pratt et al. 1998.

233 Resch et al. 1997.

234 The link was first established in 1997 as part of the Global Information Society Initiative (Global Interoperability of Broadband Networks, GIBN) supported by the G7 group of leading industrial nations. One sub-project, between Pittsburgh and Stuttgart, involved applications; a second sub-project between Sandia and Stuttgart involved distributed visualization in a virtual lab. These two projects were merged into a G-WAAT (global wide area application testbed) in the runup to Supercomputing '97. Gabriel et al. 1999, p. 131 f.

235 Metacomputing sorgt für Leistungsschub. RUS und ICA erzielen neue Weltrekorde, in: Stuttgarter Unikurier, No. 77/78, February 1998.

236 See Breckenridge 1998, pp. 20 and 22 for variants of the comet strike. The collaboration with the Albuquerque visualization group continued over the following years, but without reference to metacomputing and freed from the

zurück

collaborative constraints of the Accelerated Strategic Computing Initiative. Breckenridge et al. 2003.

237 Metacomputing sorgt für Leistungsschub. RUS und ICA erzielen neue Weltrekorde, in: Stuttgarter Unikurier, No. 77/78, February 1998.

238 HERMES had been cancelled in 1992, but developments continued in sub-projects. After 1992, developments were pursued through a variety of projects, for example, the Crew Return Vehicle and the Atmospheric Reentry Demonstrator (ARD). The ARD was used to study re-entry problems on an Ariane 5 flight in 1998. Gabriel et al. 1999, p. 139. For an easy-to-understand presentation of the problem, see Bachern and Pixius 1999, p. 56 f.

239 PSC & Stuttgart Demonstrate Latency-Free Metacomputing, in: HPCwire News Brief, 13. 11. 1998.

240 For Günter Pritschow's expectations for technology, see Pritschow and Adam 1991; Pritschow and Brandl 1997; Pritschow and Dammertz 1994; Pritschow and Daniel 1996.

241 Confirmed by Dieter Fritsch in his rector's report on the state of the university. Dieter Fritsch, Bericht zur Lage der Universität, in: Vereinigung von Freunden der Universität Stuttgart, Jahresbericht 2001, Stuttgart, p. 15 (UASt).

242 For reasons due not only to a lack of direction in the evaluation process or a lack of adaptability on the part of the organization being evaluated. Hornbostel 2016.

243 Fritsch 1982; Fritsch 1991.

244 The national supercomputing center established in Ticino in the early 1990s was a constant problem, both operationally and financially, for the Executive Board and the IT services of ETH Zurich, despite its excellent rankings. Whether a history of the Centro Svizzero e Calcolo Scientifico can ever be written without triggering every possible Helvetic sensitivity is unclear. In any case, it is no coincidence that "Manno" is only mentioned in passing in the history of ETH Zurich. Gugerli et al. 2005.

245 R. Jeltsch (et al.), Bericht: Evaluation des HLR Stuttgart, 15./16. 5. 2000, pp. 9 f. and 12 f. (Küster, private collection).

246 Ibid, p. 12 f.

247 Ibid., p. 4.

248 Ibid., p. 16.

249 Ibid., p. 3 f.

250 Wissenschaftsrat 2000.

251 Wissenschaftsrat 2002, pp. 24–28.

252 Foster/Kesselman 2011. See also Foster/Kesselman 1999 and the short survey by Resch 2006, p. 84 f.

zurück

253 This was the approach used in a project at the Leibniz Computer Center in Munich that tested interactive CFD simulations by "combining supercomputers and virtual reality." Hartlep 2003, p. 12.

254 Lindner/Resch 2003.

255 However, the Science Council proved to be a moving target, and in 2004 published its recommendation on the establishment of European supercomputers. Wissenschaftsrat 2004.

256 Hoßfeld 2003, p. 17.

257 Ibid., p. 18.

258 Ibid., p. 19.

259 Ibid., p. 23. At the suggestion of the research centers of the Helmholtz Society, the BMBF launched the German D-Grid Initiative in March 2004 with the aim of "funding projects in the fields of grid computing, e-learning, and knowledge networking with up to 100 million euros within the next five years." Neuroth et al. 2007, p. 10. The HLRS was involved in this initiative through the In-Grid project. Neuroth et al. 2007, p. 48 f.

260 Dieter Fritsch, Bericht zur Lage der Universität, in: Vereinigung von Freunden der Universität Stuttgart, Jahresbericht 2006, Stuttgart, p. 28 (UASt).

261 Ibid., p. 29.

262 In 2002/03, the HLRS underwent a "profound change," finally separating from RUS and becoming a central institution of the University of Stuttgart. The directorate changed, as did the leadership for parallel computing and visualization. From autumn 2002, former staff member Michael Resch was the new director, who returned to Stuttgart from the University of Houston, Texas, and replaced Roland Rühle, who became emeritus director in 2003. Universität Stuttgart, Rechenschaftsbericht des Rektors Professor Dr.-Ing. Dieter Fritsch, 1. 10. 2002 to 30. 9. 2003, p. 49; HLRS Annual Reports 2002, 2003, and 2004 (Resch Archive). See also the report in *Unikurier*: Schlüsselübergabe am Höchstleistungsrechenzentrum: Supercomputing zu neuen Ufern, in: Stuttgarter Unikurier, No. 91, April 2003.

263 Financing for the first phase of construction of the HLRS research building (2003–2005) is worth remarking upon. Instead of splitting the costs between the federal government and the state, as is customary, the University of Stuttgart assumed the state share through the HLRS budget. See Julia Alber, Ein Haus für Europas schnellsten Rechner, in: Stuttgarter Unikurier, No. 92, December 2003.

264 Antrag Neubau, Begründung (Fischer, Resch), 17. 8. 2000 (Resch Archive).

265 Julia Alber, Richtfest für das HLRS: Europas schnellster Rechner kann bald einziehen, in: Stuttgarter Unikurier, No. 94, December 2004.

266 For the Cray XC 40 (Hornet), which went online in 2014, the power capacity was increased from 1.8 to 5 megawatts. Landesbetrieb Vermögen und Bau Baden-Württemberg 2015, p. 41.

267 Ibid.

268 Ines Zahler, HLRS (Bundeshöchstleistungsrechenzentrum), on the website of artist Harald F. Müller, www.stratozero.net/archiv_detail.php?archivpath= archiv/2017_HLRS/ (accessed 23. 9. 2020). Vermögen und Bau Baden-Württemberg, Universitätsbauamt Stuttgart und Hohenheim, Infrastrukturerweiterung für den neuen Höchstleistungsrechner, Höchstleistungsrechenzentrum HLRS Universität Stuttgart, July 2017.

269 The Kunstverein Friedrichshafen on Müller: "The works of Harald F. Müller are invariably events and provocations in equal measure – events for seeing and provocations for thinking. Müller is not interested in categories, preferring to operate at the interstices. His works move between surface and space, image reproduction and color conception, painting and architecture. Müller is on the move, travels and lives with his eyes open to detail, making use of the overflowing well of images and color of our time with joy and relish." www.kunstverein-friedrichshafen.de/harald-f-mueller (accessed 23. 9. 2020).

270 Vermögen und Bau Baden-Württemberg, Universitätsbauamt Stuttgart und Hohenheim, Infrastrukturerweiterung für den neuen Höchstleistungsrechner, Höchstleistungsrechenzentrum HLRS Universität Stuttgart, July 2017.

271 Ibid.

272 Vermögen und Bau Baden-Württemberg, Universitätsbauamt Stuttgart und Hohenheim, Neubau HLRS Schulungszentrum, Höchstleistungsrechenzentrum HLRS, Universität Stuttgart, July 2017.

273 Ibid.

274 Ibid.

275 Andreas Kaminski, "Mit dem Zufall rechnen – Der Anspruch auf die Rettung des Glücks": description available on the Harald F. Müller website about the building of the HLRS training center, www.stratozero.net/archiv/2017_HLRS/data/hlrs-e-Kaminski-Avers%20Revers-Zwei_Seiten_Zufall-2017.pdf (accessed 23. 9. 2020).

276 The list of systems and their names is taken from the HLRS website, www.hlrs.de. It is conceivable that "Hermit" has nothing to do with a hermit person, but rather the hermit beetle and thus with the protests against the Stuttgart 21 construction project. But no written source substantiates this assertion.

zurück

277 This was also of interest to architects. See the event organized by the Ludwigsburger Architektur-Quartett on 27 October 2005 in the Ratskeller Hall in Ludwigsburg (Resch Archive). The university building office architects were invited to discuss the "merits of modern architecture" using the model of the HLRS. In addition to the HLRS, the new Ritter Museum building in Waldenbuch and the "Im Kaiser" residential development in Stuttgart were discussed.

278 Dieter Fritsch, Bericht zur Lage der Universität, in: Vereinigung von Freunden der Universität Stuttgart, Jahresbericht 2006, Stuttgart, p. 28 f. (UASt).

279 In 2003, a working group of six professors came up with a structural concept for the future development of the University of Stuttgart under the title "Zukunftsoffensive Universität Stuttgart (ZUS)." Among other things, it envisaged cancelling some humanities teacher-training courses and closing the geosciences institutes. At the time, the *Unikurier* reported "heated discussions and numerous protests." Rückblick auf eine noch nicht abgeschlossene Strukturdiskussion: Strukturell Sparen tut weh, in: Stuttgarter Unikurier, No. 92, December 2003.

280 Knie/Simon 2015, pp. 21–38. See also Leibfried 2010.

281 Universität Stuttgart, Rechenschaftsbericht des Rektors Prof. Dr.-Ing. Wolfgang Ressel, 1. 10. 2006 to 30. 9. 2007, p. 2.

282 Ibid.

283 With the participation of the HLRS, the Simulation Technology Cluster of Excellence was funded from the second round of the German federal and state governments' Excellence Initiative. According to its self-description, in the first funding phase the center "sought simulation-based solutions in six methodological areas ranging from molecular dynamics and modern mechanics to numerical mathematics and systems analysis, data management and interactive visualization, as well as high-performance computing." Universität Stuttgart, Jahresbericht 2008/09, 1. 10. 2008 to 30. 9. 2009, p. 42.

284 The Science Council's grid computing policy was geared to the German situation and no longer prioritized high-performance computing. The question of how that policy related to the shift to the European level, which interests were promoted and which opposed, is left to future research on the history of the Science Council in the 2000s. In any event, the EU launched two research initiatives during that decade that linked high-performance computing and grid computing. In the 6th Framework Programme, the Distributed European Initiative for Supercomputing Applications project (DEISA, 2004–2008) was founded with the participation of the HLRS. The project continued in the 7th Framework Programme (FP7) (DEISA2, 2008–2011). FP7 also saw the found-

zurück

ing of the BEinGRID initiative, which targeted applications in industry and business. See Gentzsch 2011a; Gentzsch 2011b.

285 Wissenschaftsrat 2014.

286 Bode et al. 2005.

287 The so-called Reuter paper (Andreas Reuter, High Performance Computing in Deutschland. Argumente zur Gründung einer strategischen Allianz. Strategiepapier 2006, Heidelberg, Resch Archive) replaced the feasibility study of 1997 (Hoßfeld et al. 1997) as an orientation tool for the German high-performance computing community. See also Protokoll des 1. Arbeitstreffens zur Koordinierung der deutschen HPC-Aktivitäten 28.3.2006, Studio der Villa Bosch, Heidelberg (Resch Archive), which discussed the "formulation of a German position with regard to EU planning in the field of HPC for FP7" and the "further steps for the development of a concept for HPC in Germany (thematic areas, working groups, expert discussions, timetable)." The minutes note that these two points were discussed together. There was "agreement on trying to develop a joint German HPC strategy similar to the BMBF contract. The prerequisite is that the BMBF contract remain intact; Mr. Reuter is instructed to clarify this as soon as possible. Incidentally, it is stated that all the elements of such a document already exist and that they only need to be compiled in suitable form."

288 Heinz-Gerd Hegering, Fortschreibung eines nat. Konzeptes. Die Rollen von PRACE, GCS und Gauss-Allianz. PowerPoint presentation, 19. 5. 2011 (Resch Archive). In 2012, the German Science Council formulated this view as follows: "In addition to a coordinated funding concept, a differentiated structure based on the division of labor for the supply of HPC resources requires close coordination between the respective operators of computing centers. These must agree on the required computer architectures and appropriate renewal periods in advance of procurement as well as jointly supported criteria for the allocation of computing time and evaluation of the scientific quality of the applications for use. At the national level, structures for this coordination among the various operators of supercomputers and high-performance computers in Germany have been created with the [GCS] and the [Gauss Allianz]." Wissenschaftsrat 2012, p. 30.

289 Heinz Gerd Hegering und Claus Axel Müller, GCS Supercomputing der Spitzenklasse. Weiterentwicklung des Gauss Centre for Supercomputing 2016–2020. PowerPoint presentation to the GCS Steering Committee, 20. 11. 2012, Slide 4 (Resch Archive).

290 Hartenstein et al. 2013. According to the annual reports, projects funded either by the Ministry of Science and the Arts (MWI) in Baden-Württemberg

or the EU predominate at the HLRS. DFG research projects, on the other hand, play only a minor role. The 2016 annual report, for example, states that 46 percent of all projects at the HLRS were funded by the EU. The state of Baden-Württemberg follows with 38 percent. See High Performance Computing Center Stuttgart (HLRS), Annual Report 2016, p. 16 f.

291 Founded in March 2008 as part of the new strategic orientation to intensify cooperation between the university and the "industrial environment," the Automotive Simulation Center Stuttgart was the first "transfer center" at the University of Stuttgart. The HLRS became the coordinating body in the center for various car company simulation projects. But it was not an independent member of the center and, thus, invisible at the decision-making level. Universität Stuttgart, Rechenschaftsbericht des Rektors Prof. Dr.-Ing. Wolfgang Ressel, 1. 10. 2006 to 30. 9. 2007, p. 3; Supercomputing für die Automobilindustrie, in: Stuttgarter Unikurier, No. 101, 2008; High Performance Computing Center Stuttgart (HLRS), Bi-Annual Report 2008/2009, p. 11.

292 Resch/Küster 2008; Resch/Wiese 2015.

293 The Stuttgart *Unikurier*, for example, reported on a contract concluded between the university and Cray in October 2010: "The agreement, which spans several years, not only includes the delivery of a new supercomputer but also covers products and services." Viele Kräne auf dem Campus Vaihingen, in: Stuttgarter Unikurier, No. 106, 2/2010, p. 11.

294 The University of Stuttgart research report cited the HLRS's many years of experience in the field of hardware and software "through cooperation with major companies (NEC, Intel, Microsoft, HP, Bull…)." Universität Stuttgart, Forschung. Entwicklung. Beratung 2008, Stuttgart 2008, p. 132 (UASt).

295 CFD and physics remained the dominant fields of application at the HLRS. While physics projects came to Stuttgart from the European region via PRACE, national projects were responsible for the consistently high proportion of CFD. The state of North Rhine-Westphalia, presumably through the Jülich Supercomputer Center, was the main user of the computers in Stuttgart outside of Baden-Württemberg. High Performance Computing Center Stuttgart (HLRS), Bi-Annual Report 2013/14, p. 29.

296 From 2000 to 2006, the courses and workshops at the HLRS averaged 237 participants per year. From 2007 to 2013 the figure rose to 522 per year, and from 2014 to 2019 to 895. PowerPoint Presentation Michael Resch, Training for HPC. Virtual Workshop: Toward a Globally Acknowledged and Free HPC Certification, 18. 5. 2020, Slide 10 (Resch Archive).

297 The CAVE (Cave Automatic Virtual Environment) was inaugurated in 2012 as a visualization center at the HLRS for 3D projects and virtual reality.

zurück

Simulierte Inneneinsichten in 3D, in: Stuttgarter Unikurier, No. 110, Issue 2, 2012, p. 11.

298 Wössner 2009. See also Currle-Linde et al. 2006.

299 Niebling et al. 2010 and Resch 2006.

300 HLRS staff doubled from roughly 70 employees in 2002 to 136 employees in 2017. High Performance Computing Center Stuttgart (HLRS), Annual Report 2002, Stuttgart 2003, p. 5; High Performance Computing Center Stuttgart (HLRS), Annual Report 2017, Stuttgart 2017, p. 92 f.

301 Wesner et al. 2010.

302 Kübert/Wesner 2010; Schubert et al. 2009.

303 Koller et al. 2015.

304 Wesner 2008.

305 Koller 2011.

306 Resch 2016.

307 Resch/Küster 2011.

308 Resch 2016. Confirmed by Michael Feldmann on the Top500 List website in July 2017 under the headline "Top500 Meanderings: Sluggish Performance Growth May Portend Slowing HPC Market," www.top500.org/news/top500-meanderings-sluggish-performance-growth-may-portend-slowing-hpc-market (accessed 23. 9. 2020).

309 Resch 2016.

310 Ibid.

311 Ibid.

312 On the question of what lay beyond Moore's law, see Cavin et al. 2012 and, specifically with reference to supercomputing, Snir 2014. At this point, referring to the imminent end of Moore's law was no longer considered heresy. See also the National Research Council's report for 2014, cited by Resch.

Bibliography

Archival materials from the Stuttgart University Archive, various records of the HLRS, and documents made available to us by Uwe Küster, Jutta Sauer, Erika Fischer, and Michael Resch are quoted in detail in the footnotes. All publications cited by author and year are listed below.

Ahlrichs, Reinhart 1993: Parallelrechner versus Workstation Cluster. Positionspapier, in: Hans-Werner Meuer (ed.): *Supercomputer '93. Anwendungen, Architekturen, Trends. Mannheim 24.–26. Juni 1993. Proceedings*, Berlin etc., pp. 179–180.

Anthes, Gary H. 1994: Research Lab Sizes Up Slew of Supercomputers, in: *Computerworld*, 28 (31), p. 55.

Bachern, Achim, Kay Pixius 1999: Simulation & Mathematik: Anwendungen in der Luft- und Raumfahrt und in der Verkehrsforschung, in: Andreas Dress, Gottfried Jäger (eds.): *Visualisierung in Mathematik, Technik und Kunst*, Braunschweig, Wiesbaden, pp. 51–62.

Bailey, David H. 1992: Misleading Performance Reporting in the Supercomputing Field, Technical Report RNR-92-005. Dec 1, 1992, in: *Scientific Programming*, 1 (2), pp. 141–151.

Beisel, Thomas et al. 1997: *An Extension to MPI for Distributed Computing on MPPs*, Berlin, Heidelberg, pp. 75–82.

Bode, Arndt et al. 2005: *Petaflop-Computing mit Standort Deutschland im europäischen Forschungsraum, Bedarf und Perspektiven aus Sicht der computergestützten Natur- und Ingenieurwissenschaft. "Scientific Case" im Auftrag des BMBF*, Bonn.

Brand, Wolfgang 2012: Argyris und die Frühgeschichte des Höchstleistungsrechnens, in: Klaus Hentschel (ed.): *Zur Geschichte von Forschungstechnologien. Generizität – Interstilialität – Transfer*, Diepholz, Berlin, Stuttgart. Verlag für Geschichte der Naturwissenschaften und der Technik, pp. 206–220.

Breckenridge, Arthurine R. 1998: Teraflop Visualization, in: Martin Göbel et al. (eds.): *Virtual Environments '98*, Vienna, pp. 18–27.

Breckenridge, Arthurine R. et al. 2003: Distributed, On-Demand, Data-Intensive and Collaborative Simulation Analysis, in: *Future Generation Computer Systems*, 19 (6), pp. 849–859.

zurück

Brune, Matthias A. et al. 1999: Message-Passing Environments for Metacomputing, in: *Future Generation Computer Systems*, 15 (5–6), pp. 699–712.

Cale, Edward Graham et al. 1979: Price/Performance Patterns of U.S. Computer Systems, in: *Communications of the ACM*, 22 (4), pp. 225–233.

Cavin, Ralph K. et al. 2012: Science and Engineering beyond Moore's Law, in: *Proceedings of the IEEE*, 100, pp. 1720–1749.

Commisson of the European Countries. 1992: Report of the EEC Working Group on High-Performance Computing. Rubbia Report (Reprint of the Original Report, Published in February 1991), in: Hans-Werner Meuer: *Supercomputer '92. Anwendungen, Architekturen, Trends, Seminar, Mannheim, 25.–27. Juni 1992*, Berlin etc., pp. 213–232.

Currle-Linde, Natalia et al. 2006: The Role of Supercomputing in Industrial Combustion Modeling, in: Thomas Bönisch et al. (eds.): *High Performance Computing on Vector Systems. Proceedings of the High Performance Computing Center Stuttgart, March 2005*, Berlin, Heidelberg, pp. 109–122.

Deutsche Forschungsgemeinschaft 1991: *Kommission für Rechenanlagen. Zur Ausstattung der Hochschulen in der Bundesrepublik Deutschland mit Datenverarbeitungskapazität für die Jahre 1992 bis 1995*, Bonn.

Doltsinis, Ioannis 2004: Obituary John Argyris, in: *Communications in Numerical Methods in Engineering*, 20, pp. 665–669.

Dongarra, Jack J. et al. 1979: *LINPACK Users' Guide*, Philadelphia.

Dongarra, Jack J., Piotr Luszczek, Antoine Petite 2003: The LINPACK Benchmark: Past, Present and Future, in: *Concurrency and Computation: Practice and Experience* 15, pp. 803–810.

Dongarra, Jack J., Piotr Luszczek 2011: Top500, in: David Padua (ed.): *Encyclopedia of Parallel Computing*, Boston, MA, pp. 2055–2057.

Effenberger, Franz 2020: *Lothar Späths Forschungsförderung und Technologiepolitik am Beispiel der Universität Stuttgart*, Ubstadt-Weiher.

European Commission 1994: *Growth, Competitiveness, and Employment. White Paper Follow-Up Report on Europe and the Global Information Society. Interim Report on Trans-European Networks. Progress Report on Employment. Extracts of the Conclusions of the Presidency of the Corfu European Council* (Bulletin of the European Union, Supplement 2/94), Brussels and Luxemburg (Bangemann Report).

Evers, Adalbert, Helga Nowotny 1987: *Über den Umgang mit Unsicherheit. Die Entdeckung der Gestaltbarkeit von Gesellschaft*, Frankfurt am Main.

Foster, Ian, Carl Kesselman (eds.) 1999: *The Grid: Blueprint for a New Computing Infrastructure*, San Francisco.

zurück

Foster, Ian, Carl Kesselman 2011: The History of the Grid, in: Ian Foster, Carl Kesselman (eds.): *High Performance Computing: From Grids and Clouds to Exascale*, Amsterdam, pp. 3–30.

Fritsch, Dieter 1982: *Entwurf digitaler zweidimensionaler nichtrekursiver Filter*, Munich.

Fritsch, Dieter 1991: *Raumbezogene Informationssysteme und digitale Geländemodelle*, Munich.

Gabriel, Edgar et al. 1997: *Erweiterung einer MPI-Umgebung zur Interoperabilität verteilter MPP-Systeme*, Studienarbeit Angewandte Informatik, Anwendungen der Informatik im Maschinenwesen (RUS-37), Stuttgart.

Gabriel, Edgar et al. 1999: High Performance Metacomputing in a Transatlantic Wide Area Application Testbed, in: Gene Cooperman et al. (eds.): *Workshop on Wide Area Networks and High Performance Computing* (Lecture Notes in Control and Information Sciences, 249), London, pp. 131–142.

Gall, Alexander 2001: "Bundesliga-Spielregeln in der Wissenschaftspolitik." Föderalismus und Forschungspolitik zur Mikroelektronik, in: Johannes Abele et al. (eds.): *Innovationskulturen und Fortschrittserwartungen im geteilten Deutschland*, Köln, pp. 147–164.

Geiger, Alfred 1993: PARIS – Das Parallel RISe Projekt der IBM, DLR und Universität Stuttgart, in: Hans-Werner Meuer (eds.): *Supercomputer '93. Anwendungen, Architekturen, Trends*, Berlin, Heidelberg, pp. 166–178.

Geiger, Alfred 1994: *Konzeption, Stand und Betrieb eines bundesweit zugänglichen Höchstleistungsrechnersystems an der Universität Stuttgart (Folienset), zK&I* (Zentren für Kommunikation und Informationsverarbeitung in Lehre und Forschung e.V.), Mannheim.

Geist, Al 1994: PVM: Parallel Virtual Machine – A Users' Guide and Tutorial for Networked Parallel Computing, Scientific and Engineering Computation, Cambridge, MA.

Gentzsch, Wolfgang 1987: Vergleich der Vektorrechner CRAY-2, CRAY X-MP, Fujitsu VP-200, IBM 3090-VF und UNISYS-ISP, in: *Praxis der Informationsverarbeitung und Kommunikation*, 10 (2), pp. 93–102.

Gentzsch, Wolfgang 2011a: DEISA – Distributed European Infrastructure for Supercomputing Applications, in: *Journal of Grid Computing*, 9, pp. 255–277.

Gentzsch, Wolfgang 2011b: Porting HPC Applications to Grid and Clouds, in: Emmanuel Udoh (ed.): *Cloud, Grid and High Performance Computing. Emerging Applications*, Hershey, PA, pp. 10–38.

Gropp, William et al. 1994: Using MPI: Portable Parallel Programming with the Message-Passing Interface, Cambridge, MA.

zurück

Gugerli, David 2018: *Wie die Welt in den Computer kam. Zur Entstehung digitaler Wirklichkeit*, Frankfurt am Main.

Gugerli, David, Magaly Tornay 2018: Das Zeitalter der Konfigurationen 1980 bis 2010. Ein Beitrag zur zeithistorischen Debatte, in: *Historische Anthropologie*, 26 (2), pp. 224–244.

Gugerli, David et al. 2005: *Transforming the Future: ETH Zurich and the Construction of Modern Switzerland 1855–2005*, Zurich.

Harms, Uwe, Hans-Werner Meuer 1995: Höchstleistungsrechnen in Deutschland – ein Rückblick, in: *Praxis der Informationsverarbeitung und Kommunikation*, 18 (2), pp. 100–107.

Hartenstein, Hannes et al. 2013: Aktuelle Umsetzungskonzepte der Universitäten des Landes Baden-Württemberg für Hochleistungsrechnen und datenintensive Dienste, in: *Praxis der Informationsverarbeitung und Kommunikation*, 36 (2), pp. 99–108.

Hartlep, Thomas 2003: Interactive CFD Simulations by Coupling Supercomputers with Virtual Reality, in: *InSiDE*, 1 (2), pp. 12–13.

Hertweck, Friedrich et al. 1979: *Benchmark-Versuche mit der Cray-1*, Garching.

Hoffmann, Christoph 2004: Eine Maschine und ihr Betrieb: Zur Gründung des Recheninstituts der Technischen Hochschule Stuttgart, 1956–1964, in: Barbara Büscher et al. (eds.): *Ästhetik als Programm: Max Bense/Daten und Streuungen*, Berlin, pp. 118–129.

Hord, R. Michael 1982: *The Illiac IV. The First Supercomputer*, Berlin, Heidelberg, New York.

Hornbostel, Stefan 2016: (Forschungs-)Evaluation, in: Dagmar Simon et al. (eds.): *Handbuch Wissenschaftspolitik*, Wiesbaden, pp. 243–260.

Hoßfeld, Friedel 1984: Vektorrechner, in: *Physikalische Blätter*, 40 (8), pp. 280–281.

Hoßfeld, Friedel et al. 1997: *Verbund der Supercomputer-Zentren in Deutschland – eine Machbarkeitsanalyse*, Jülich.

Hoßfeld, Friedel 1998: *Verbund der Supercomputer-Zentren in Deutschland – Ansichten, Einsichten, Aussichten*, Jülich.

Hoßfeld, Friedel 2003: National Supercomputing Centres – Turning the Innovation Spiral, in: *InSiDE*, 1 (2), pp. 16–23.

HPSC, Arbeitsgruppe 1992: *Memorandum zur Initiative High Performance Scientific Computing (HPSC): Situation und Erfordernisse des wissenschaftlichen Höchstleistungsrechnens in Deutschland*. Positionspapiere der Arbeitsgruppen, n.p.

Jenkins, Richard A. 1986: *Supercomputers of Today and Tomorrow: The Parallel Processing Revolution*, Blueridge Summit, PA.

zurück

Khokhar, Ashfaq Ahmad 1993: Heterogeneous Computing: Challenges and Opportunities, in: *IEEE Computer*, 26 (6), pp. 18–27.

Knie, Andreas, Dagmar Simon 2016: Innovation und Exzellenz: Neue und alte Herausforderungen für das deutsche Wissenschaftssystem, in: Dagmar Simon et al. (eds.): *Handbuch Wissenschaftspolitik*, Berlin, pp. 21–38.

Koller, Bastian 2011: *Enhanced SLA Management in the High Performance Computing Domain (HLRS-08)*, University of Stuttgart.

Koller, Bastian et al. 2015: Towards an Environment to Deliver High Performance Computing to Small and Medium Enterprises, in: Michael Resch et al. (eds.): *Sustained Simulation Performance 2015. Proceedings of the Joint Workshop on Sustained Simulation Performance, University of Stuttgart (HLRS) and Tohoku University, 2015*, Cham, pp. 41–50.

Kolsky, H. G. 1968: Computer in der Meteorologie / Some computer aspects of meteorology, in: *it – Information Technology*, 10 (1), pp. 265–279.

Kübert, Roland, Stefan Wesner 2010: Service Level Agreements for Job Control in High-Performance Computing, in: Maria Ganzha, Marcin Paprczyki (eds.): *Proceedings of the International Multiconference on Computer Science and Information Technology, October 18–20, 2010, Wisła, Poland*, 5, Katowicze, Los Alamitos, CA, pp. 655–661.

Lang, Ulrich 1989: Modularisierungskonzepte bei Simulationsproblemen in integrierten Programmsystemen, von der Universität Stuttgart zur Erlangung der Würde eines Doktor-Ingenieurs (Dr.-Ing.) genehmigte Abhandlung (RUS-2), Stuttgart.

Lang, Ulrich 1993: Graphische Basissoftware auf Höchstleistungsrechnern, in: Hans-Werner Meuer (ed.): *Basissoftware für Vektor- und Parallelrechner*, Munich etc., pp. 126–138.

Lang, Ulrich, Roland Rühle 1992: Visualisierung von Supercomputerberechnungen am netzintegrierten Ingenieursarbeitsplatz, in: Hans-Werner Meuer (ed.): *Heterogene Netze und Supercomputer*, Berlin, pp. 121–133.

Lax, Peter D. 1982: Report of the Panel on Large Scale Computing in Science and Engineering, www.pnnl.gov/scales/docs/lax_report1982.pdf.

Leibfried, Stephan (ed.) 2010: *Die Exzellenzinitiative, Zwischenbilanz und Perspektiven*, Frankfurt am Main, New York.

Leimbach, Timo 2011: *Die Softwarebranche in Deutschland. Entwicklung eines Innovationssystems zwischen Forschung, Markt, Anwendung und Politik von 1950 bis heute*, Stuttgart.

Lindner, Peggy, Michael Resch 2003: Gridwelten. User Requirements and Environments for Grid Computing, in: *DFN Mitteilungen*, 62 (6), pp. 13–15.

zurück

Loebich, Isabel et al. 1987: DFN-RSYST. Software zur Unterstützung verteilter Anwendungen in heterogenen Rechnernetzen, in: Nina Gerner, Otto Spaniol (eds.): *Kommunikation in verteilten Systemen*, Berlin, Heidelberg, pp. 653–664.

MacKenzie, Donald 1991: The Influence of the Los Alamos and Livermore National Laboratories on the Development of Supercomputing, in: *Annals of the History of Computing*, 13 (2), pp. 179–201.

Meuer, Hans-Werner 1990: Vorwort, in: Hans-Werner Meuer (ed.): *Supercomputer '90. Anwendungen, Architekturen, Trends. Mannheim 21.–23. Juni 1990. Proceedings*, Berlin etc., pp. V–VI.

Meuer, Hans-Werner (ed.) 1991a: *Parallelisierung komplexer Probleme. Einsatz von Parallelrechnern in Forschung und Industrie*, Berlin.

Meuer, Hans-Werner 1991b: Vorwort, in: Hans-Werner Meuer (ed.): *Supercomputer '91. Anwendungen, Architekturen, Trends. Mannheim 20.–22. Juni 1991. Proceedings*, Berlin etc., pp. V–VI.

Meuer, Hans-Werner 1992: Parallelrechner – bringen die 90er Jahre den Durchbruch? (erschienen in *Praxis der Informationsverarbeitungn und Kommunikation*, 1, 1990), in: Hans-Werner Meuer (ed.): *Supercomputer 1986–1990. Anwendungen, Architekturen, Trends*, Munich etc., pp. 387–393.

Meuer, Hans-Werner, Erich Strohmaier 1996: Supercomputing auf dem Weg ins Jahr 2000, in: Hans-Werner Meuer (ed.): *Supercomputer 1996. Anwendungen, Architekturen, Trends*, Munich etc., pp. 181–192.

Ministerium für Wissenschaft und Forschung Baden-Württemberg 1995: *Informations- und Kommunikationstechnik in den Hochschulen des Landes Baden-Württemberg, Ausstattungsplan für den Zeitraum 1995 bis zum Jahr 2000 (EDV-Gesamtplan IV)*, Stuttgart.

Ministerium für Wissenschaft und Kunst Baden-Württemberg 1980: *Fortschreibung des Gesamtplans zur Ausstattung der Hochschulen des Landes Baden-Württemberg mit Datenverarbeitungsanlagen für den Zeitraum 1979–1985 (EDV Gesamtplan II)*, Stuttgart.

Ministerium für Wissenschaft und Kunst Baden-Württemberg 1989a: *Gesamtplan zur Ausstattung der Hochschulen des Landes Baden-Württemberg mit Datenverarbeitungsanlagen für den Zeitraum 1989–1994 (EDV-Gesamtplan III)*, Stuttgart.

Ministerium für Wissenschaft und Kunst Baden-Württemberg 1989b: *Kommission Forschung. Baden-Württemberg 2000*. Abschlussbericht, Stuttgart.

Ministerium für Wissenschaft und Kunst Baden-Württemberg, Kommission Forschung 1989: *Abschlussbericht*, Stuttgart.

Mönkediek, Joachim 2009: Rückblick auf die Förderung der Informationstechnologie an Hochschulen durch Bund und Länder, in: Hannes Hartenstein (ed.): *Informationstechnologie und ihr Management im Wissenschaftsbereich. Festschrift für Prof. Dr. Wilfried Juling*, Karlsruhe, pp. 23–28.

National Research Council 2014: *Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017–2020: Interim Report*, Washington, DC.

Neuroth, Heike et al. (eds.) 2007: *Die D-Grid Initiative*, Göttingen.

Niebling, Florian et al. 2010: Collaborative Steering in Post-Processing of Simulation on HPC Resources: Everyone, Anytime, Anywhere, in: Marcelo Knorich Zuffo et al. (eds.): *Web3D '10. Proceedings of the 15th International Conference on Web 3D Technology*, New York, pp. 101–108.

Pias, Claus 2011: On the Epistemology of Computer Simulation, in: *Zeitschrift für Medien- und Kulturforschung*, 1, pp. 29–54.

Pratt, Thomas J. et al. 1998: *Sandia's Network for SC '97: Supporting Visualization, Distributed Cluster Computing, and Production Data Networking with a Wide Area High Performance Parallel Asynchronous Transfer Mode (ATM) Network* (Sandia Report SAND98-1154), Albuquerque.

Pritschow, Günter, Wolfgang Adam 1991: *Leit- und Steuerungstechnik in flexiblen Produktionsanlagen* (Fortschritte der Fertigung auf Werkzeugmaschinen), Munich.

Pritschow, Günter, Thomas Brandl 1997: *Schnittstellen im CAD CAM-Bereich* (Fortschritte der Fertigung auf Werkzeugmaschinen), Munich.

Pritschow, Günter, Ralph Dammertz 1994: *Roboteranwendung für die flexible Fertigung* (Fortschritte der Fertigung auf Werkzeugmaschinen), Munich.

Pritschow, Günter, Christian Daniel 1996: *Komponenten und Schnittstellen für offene Steuerungssysteme* (Fortschritte der Fertigung auf Werkzeugmaschinen), Munich.

Reinsch, Karl-Gottfried 1980: Regionale Großrechenzentren. Zur Versorgung Japanischer Universitäten, in: *Das Rechenzentrum*, 3 (2), pp. 73–79.

Reinsch, Karl-Gottfried 1982: Strukturveränderungen des Rechenzentrums und seiner Benutzer, in: *Das Rechenzentrum*, 31, pp. 171–175.

Reinsch, Karl-Gottfried 1987: Die Cray-2 der Universität Stuttgart. Eine Herausforderung an Wissenschaft und Industrie, in: *Praxis der Informationsverarbeitung und Kommunikation*, 10 (2), pp. 103–108.

Reinsch, Karl-Gottfried 1988: Der Supercomputer und die wissenschaftliche Erkenntnis, in: *Helvetica Physica Acta*, 61 (1–2), pp. 9–23.

zurück

Resch, Michael 2006: Grid at the interface of industry and research, in: Thomas Barth, Anke Schüll (eds.): *Grid Computing: Konzepte – Technologien – Anwendungen*, Wiesbaden, pp. 85–98.

Resch, Michael 2016: Perspectives in High Performance Computing, in: Vinai K. Singh et al. (eds.), *Modern Mathematical Methods and High Performance Computing in Science and Technology*, Singapore, pp. 169–178.

Resch, Michael, Uwe Küster 2008: HPC in Industrial Environments, in: Egon Krause et al. (eds.): *Computational Science and High Performance Computing III. The 3rd Russian-German Advanced Research Workshop*, Berlin, Heidelberg, pp. 8–13.

Resch, Michael M., Uwe Küster 2011: HPC Processor Technologies and Their Impact on Simulation, in: Egon Krause (ed.): *Computational Science and High Performance Computing IV. Notes on Numerical Fluid Mechanics,* 115, Berlin, Heidelberg, pp. 17–28.

Resch, Michael, Andreas Wiese 2015: Industrial Applications on Supercomputing in Germany, in: Anwar Osseyran, Merle Giles (eds.): *Industrial Applications of High Performance Computing. Best Global Practices*, Boca Raton, FL, pp. 63–78.

Resch, Michael et al. 1994: Message-Passing-Systems on Workstation Clusters and Parallel Computers – the Impact of Software and Network Architectures on Applications, in: Wolfgang Gentzsch, Uwe Harms (eds.): *High-Performance Computing and Networking. International Conference and Exhibition Munich, Germany, April 18–20, 1994 Proceedings,* 2: *Networking and Tools*, New York, pp. 260–266.

Resch, Michael et al. 1997: PACX-MPI, in: *BI. Informationen für Nutzer des Res chenzentrums*, (11/12), pp. 13–14.

Reuter, Andreas 1993: Parallelrechner versus Workstation Cluster. Positionspapier, in: Hans-Werner Meuer (ed.): *Supercomputer '93. Anwendungen, Architekturen, Trends. Mannheim 24.–26. Juni 1993. Proceedings*, Berlin, pp. 181–183.

Rühle, Roland 1971: Datenstruktur und Steuerung des Reaktorprogrammsystems RSYST, in: Deutsches Atomforum (ed.): *Reaktortagung 1971*, Bonn, pp. 38–41.

Rühle, Roland 1973: *RSYST, ein integriertes Modulsystem mit Datenbasis zur Automatisierten Berechnung von Kernreaktoren*, Stuttgart.

Rühle, Roland 1990: *RSYST. Ein Softwaresystem zur Integration von Daten und Programmen zur Simulation wissenschaftlich-technischer Systeme*, Stuttgart.

zurück

Rühle, Roland 1995: Deutsches Supercomputerzentrum?, in: Hans-Werner Meuer (ed.): *Supercomputer 1995. Anwendungen, Architekturen, Trends*, Munich etc., pp. 262–264.

Sandner, Peter 2008: ALWR-BW. Arbeitskreis der Leiter der Universitätsrechenzentren in Baden-Württemberg, in: *Praxis der Informationsverarbeitung und Kommunikation*, 31 (3), pp. 193–198.

Schlecht, Bernhard 1989: *Entwurf und Implementierung eines interaktiven anwendungsunabhängigen Grafiksystems*, von der Universität Stuttgart zur Erlangung der Würde eines Doktor-Ingenieurs (Dr.-Ing.) genehmigte Abhandlung (RUS-3), Stuttgart.

Schubert, Lutz et al. 2009: Service Oriented Operating Systems: Future Workspaces, in: *IEEE Wireless Communication*, 16 (3), pp. 42–50.

Simon, Horst D. 1995: Parallel Supercomputing 1987–2000, in: Hans-Werner Meuer (ed.): *Supercomputer 1995. Anwendungen, Architekturen, Trends*, Munich etc., pp. 22–30.

Smarr, Larry, Charles E. Catlett 1992: Metacomputing, in: *Communications of the ACM*, 35 (6), pp. 44–52.

Snir, Marc 2014: *The Future of Supercomputing, Proceedings of the 28th ACM International Conference on Supercomputing*, Munich, pp. 261–262.

Späth, Lothar 1985: *Wende in die Zukunft die Bundesrepublik auf dem Weg in die Informationsgesellschaft*, Spiegel-Buch, Reinbek bei Hamburg.

Trottenberg, Ulrich 2008: Das Superrechnerprojekt SUPRENUM (1984–1989), in: Bernd Reuse, Roland Vollmer (eds.): *Informatikforschung in Deutschland*, Berlin, Heidelberg, New York, pp. 176–187.

Übelmesser, Robert 1989: Die neuen Supercomputer von Cray, in: Hans-Werner Meuer (ed.): *Supercomputer '89. Anwendungen, Architekturen, Trends*, Berlin etc., pp. 31–42.

Universität Stuttgart 2003: *Simulation und Visualisierung*, Stuttgart.

Wagner, Siegfried 2009: German and EU High-Performance Computation Centers, in: Ernst Heinrich Hirschel, Egon Krause (eds.): *100 Volumes of "Notes on Numerical Fluid Mechanics." 40 Years of Numerical Fluid Mechanics and Aerodynamics in Retrospect*, Berlin, Heidelberg, pp. 45–57.

Weingart, Peter 2006: *Das Wissensministerium. Ein halbes Jahrhundert Forschungs- und Bildungspolitik in Deutschland*, Weilerswist.

Weisman, Jonathan 1995: New Mission for the National Labs, in: *Science*, 270 (5233), pp. 20–22.

Wesner, Stefan 2008: *Integrated Management Framework for Dynamic Virtual Organisations* (HLRS-06), Stuttgart.

zurück

Wesner, Stefan et al. 2010: Challenges and Opportunities of Hybrid Computing Systems, in: Sabine Roller et al. (eds.): *High Performance Computing on Vector Systems 2009*, Berlin, Heidelberg, pp. 31–39.

Wiegand, Josef 1994: *Informatik und Großforschung. Geschichte der Gesellschaft für Mathematik und Datenverarbeitung* (Studien zur Geschichte der deutschen Großforschungseinrichtungen), Frankfurt am Main, New York.

Wissenschaftsrat 1985: *Wettbewerbsempfehlung des WR. Empfehlungen zum Wettbewerb im deutschen Hochschulsystem*, Köln.

Wissenschaftsrat 1987: *Empfehlung zur Ausstattung der Hochschulen mit Rechenkapazität*, Köln.

Wissenschaftsrat 1995a: *Empfehlung zur Versorgung von Wissenschaft und Forschung mit Höchstleistungsrechenkapazität*, Kiel.

Wissenschaftsrat 1995b: *Empfehlungen zur Bereitstellung leistungsfähiger Kommunikationsnetze für die Wissenschaft*, Saarbrücken.

Wissenschaftsrat 2000: *Empfehlung zur künftigen Nutzung von Höchstleistungsrechnern*, Mainz.

Wissenschaftsrat 2002: *Stellungnahme zur Anmeldung des Landes Baden-Württemberg auf Errichtung eines Höchstleistungsrechner-Systems in Karlsruhe und Stuttgart zum 32. Rahmenplan*, Mannheim.

Wissenschaftsrat 2004: *Empfehlung zur Einrichtung europäischer Höchstleistungsrechner*, Hamburg.

Wissenschaftsrat 2012: *Strategische Weiterentwicklung des Hoch- und Höchstleistungsrechnens in Deutschland*. Positionspapier, Berlin.

Wissenschaftsrat 2014: *Bedeutung und Weiterentwicklung von Simulation in der Wissenschaft*. Positionspapier, Dresden.

Wössner, Uwe 2009: *Virtuelle und hybride Prototypen in kooperativen Arbeitsumgebungen* (HLRS-07), Stuttgart.

List of Figures

zurück