SWISS FEDERAL INSTITUTE OF TECHNOLOGY
ECOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

N° 10 – NOV. 98

EPFL SUPERCOMPUTING REVIEW

NUMERICAL SIMULATION FOR SCIENCE AND TECHNOLOGY

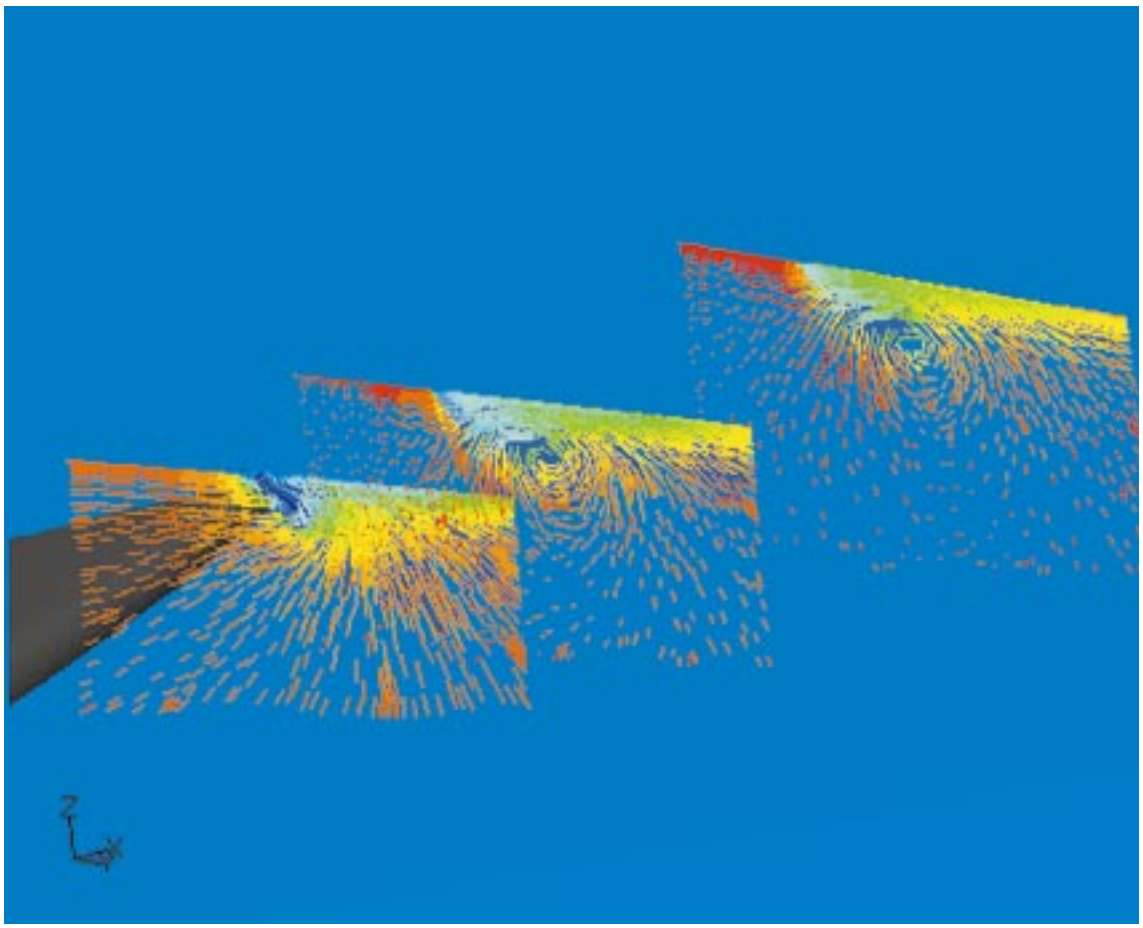

Vortex separation behind the transom of an IACC hull computed by FLUENT/UNS
(see article on page 24)

ECOLE POLYTECHNIQUE
FEDERALE DE LAUSANNE

CONTENTS
TABLE DES MATIÈRES

Parallel implementation of graph diameter algorithms

Implémentation parallèle d'algorithmes pour le diamètre des graphes

Jean-Albert Ferrez, Komei Fukuda, & Thomas M. Liebling 3

An information system for the integrated cost management in the constrcution industry

Un système d'information pour la gestion intégrée des coûts dans la construction

Abdelilah Zertiti 6

A hierarchical computer simulation model for the evolution of the microstructure produced by displacement cascades in metals

Un modèle hiérarchique de simulation numérique pour l'évolution de la microstructure produite par irradiation aux ions lourds dans les métaux

Abderrahim Almazouzi , Maria José Caturla, Tomas Diaz de la Rubia & Maximo Victoria 10

Swiss-Tx: first experiences on the T0 system

Swiss-Tx: premières expériences sur le système T0

Ralf Gruber & Yves Dubois Pèlerin 19

America's Cup Yacht Design using Advanced Numerical Flow Simulations

Conception d'un yacht de classe America en utilisant des méthodes de pointe pour simuler les écoulements hydro- et aérodynamiques

Bernard Bonjour, Mario Caponnetto, Alessandro Castelli, Pierre-Louis Mathey, Stephane Sanchi & Mark L. Sawley 24

Pour en savoir plus sur CAPA: http://capawww.epfl.ch/

EDITORIAL

A l'image des activités de notre Ecole, le bouquet de la cuvée 98 est fait de pluridisciplinarité, ouverture et esprit pionnier: mathématiques, physique et sciences des matériaux, hydro et aérodynamique, gestion des flux d'informations: quatre applications des sciences de l'ingénieur nécessitant d'importantes ressources informatiques sont ici présentées. Le cinquième article relate l'exploration d'une nouvelle solution pour le calcul de puissance. Toutes ces contributions visent à augmenter l'efficacité, que ce soit par une meilleure compréhension des phénomènes, par une meilleure gestion ou en anticipant les changements dans le monde environnant.

De nombreuses années après que certains pionniers dans cette Ecole se soient lancés la voie de l'informatique de puissance, l'expérience et le savoir faire des ingénieurs et scientifiques constituent dans ce domaine aujourd'hui la principale plus-value. Les actions que des groupes d'intérêts, tels que CAPA, pourraient choisir de mener sur le long terme – consolidation du leadership en applications scientifiques, valorisation par la formation, promotion des applications et méthodes auprès du monde industriel, des celui des affaires et de la finance – s'appuieront largement sur cette compétence.

Réaliser un tel journal eut été impossible sans une collaboration sans faille avec les auteurs, la compétence et la patience de l'éditrice technique: qu'ils soient ici tous remerciés.

Marie-Christine Sawley

EDITORIAL

Aligned with the activities at the EPFL, the 98 vintage of the Supercomputing Review is made of multidisciplinarity, with a pioneer and open character: mathematics, material science, hydro and aerodynamics, information management, four engineering applications requiring important computer resources are presented. How to serve the needs of the scientific community by developping a totally new concept in high performance computing is the subject of the fifth article. Their common goal is to increase efficiency, either by giving a deeper understanding of the natural phenomena, supplying better management or by anticipating the changes in the surrounding world.

Many years have gone by since the pioneering times at the EPFL in computer intensive applications. Today, the skills and competences of its engineers and scientists represent its strongest asset in this field. Groups like CAPA may choose long term actions based on such a platform: keeping and developping a leader role in engineering and scientific applications, training and education, promoting and transferring applications and methods into the business, finance and industrial worlds.

Producing such a journal would have been impossible without a faultless collaboration with the authors, competence and patience shown by the technical editor: my personal acknowledgements to all of them.

Marie-Christine Sawley

PARALLEL IMPLEMENTATION OF GRAPH DIAMETER ALGORITHMS

JEAN-ALBERT FERREZ, KOMEI FUKUDA, THOMAS M. LIEBLING, EPFL, DÉPARTEMENT DE MATHÉMATIQUES

Le diamètre d'un graphe est la longueur maximale d'un plus court chemin entre deux sommets du graphe. Les méthodes pour le calculer peuvent être rangées en deux catégories : celles qui manipulent la matrice des distances et celles qui appliquent plusieurs fois un algorithme de plus courts chemins à source unique tel que celui de Dijkstra. Nous présentons plusieurs implémentations de ces dernières sur les ordinateurs parallèles de l'EPFL et discutons leurs performances, principalement sur de grands graphes peu denses.

The diameter of a graph is the maximum length of shortest paths between two vertices in the graph. The methods to compute it can be classified in two categories: manipulation of the distance matrix and repeated use of a single source shortest path algorithm like Dijkstra's. We present several parallel implementations of the latter on the parallel machines available at EPFL, and discuss their performances, focusing on large, sparse, polytopal graphs.

INTRODUCTION

Let G be a graph with N vertices and M edges. The diameter of G is defined as the maximum length of shortest paths between two vertices of G. This value, together with other characteristics, gives a hint on the general shape or structure of the graph. Several algorithms use it as a lower or upper bound. In the context of parallel computing, the diameter of the network interconnecting the nodes of a distributed system is a measure of the worst-case communication time. For many regular graphs like rings, hypercubes, torus... the diameter is a function of the number of nodes, but in general, it can take any value and is not known in advance.

THE TEST GRAPHS

We tested the implementations on some regular graphs (rings and hypercubes) and on some polytopal graphs. Those graphs are the wire-frame representation of a polytope in $\mathbf{R}^d$. The set of vertices is the same, the set of edges is the same. In 1957 Hirsch conjectured that the maximal diameter of the graphs of polyhedra in $\mathbf{R}^d$ with f facets $\Delta(d,f)$ obeys: $\Delta(d,f) \leq f - d$. Klee and Walkup proved it is false for unbounded polyhedra, but the conjecture is still open for polytopes [2]. Nowadays, the availability of high performance computing resources allows checking the validity of the conjecture on a range of much larger graphs.

We used the graphs of dual-cyclic polytopes of dimension d with 2d facets, for d = 10,11,12 and 13 (Table 1). The dual-cyclic polytopes attain the maximum number of vertices for a fixed dimension and a fixed number of facets. Their diameters are not known in general but our codes verified that the Hirsh upper bound d = 2d - d is tight for these cases [7].

Dimension	Number of vertices	Number of edges	Approx. memory footprint (Mb)
10	4004	20020	0.5
11	8736	48048	1
12	24752	148512	3
13	54264	352716	9

Table 1 – The test graphs

In this context, the actual length of an edge is not important, as the length of a path is given by the number of edges traversed, not by the sum of their length. This is why we focus in this work on undirected, unweighted graphs. If some non-negligible changes are required to support directed graphs, all the codes are designed to treat both weighted and unweighted graphs in the most efficient way.

THE MATRIX METHODS

Since the diameter is the greatest distance between any two vertices, we could simply compute those N^2 (in fact N(N-1)/2) distances and take the maximum. This can be done by manipulating an NxN matrix: the idea is to keep for every pair of vertices (i,j) either a value or a special mark that says "I do not know any path of that kind between i and j'', and then to have these values or marks evolve towards the length of the shortest path between i and j. Ultimately, the greatest element in this matrix is the diameter. These methods are very simple, but - unfortunately - not efficient. They require the storage of the whole NxN distance matrix, which is not possible for large graphs. The best methods of this category perform $O(N^3)$ operations, which could result in years of CPU time! Last but not least, when considering a parallel implementation, they suffer from a poor computation / communication ratio and show

almost no data-locality. In practice, they are interesting for small ($N<100$), dense graphs ($M \approx N^2$), but useless for sparse graphs with more than a thousand vertices. They are however used in tutorials and theoretical studies [4,8]. Although we have implementations for the Cray T3D and the SGI Origin 2000, we won't address them here. Issues such as data placement and consistency, one-side communication, synchronization and scalability are covered in [3].

THE NON-MATRIX METHODS

Instead of explicitly computing and storing the NxN distances we can consider only the distances between selected pairs in the graph. Since the length of every shortest path in G is a lower bound for the diameter, repeatedly applying a single source shortest paths algorithm - and remembering the current greatest shortest path - is in fact improving this lower bound until it reaches the value of the diameter. This happens either when all sources have been processed, or when this lower bound meets some other upper bound on the diameter. Such an upper bound is provided by the diameter of the shortest path spanning tree returned by Dijkstra's algorithm, which is easy to compute with a simple *hanging* (fig. 1). More details on Dijkstra's algorithm, its implementation and the bounding technique can be found in [3,6]. In the worst case, that is when the two bounds never meet and all the sources have to be computed, the method still performs between $O(N^2)$ and $O(N^3)$ operations and only a major theoretical breakthrough could improve this.

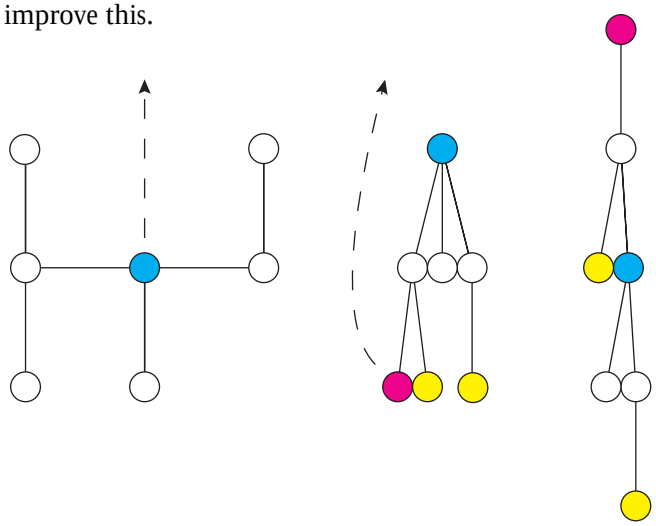

Fig. 1 – Hanging a tree twice to compute its diameter. The shortest paths spanning tree is already hung by its root.

The non-matrix methods have some nice properties when it comes to their parallel implementation:

Local Dijkstra

Although parallel implementations of Dijkstra's algorithm exist [1], they rely on a distributed priority queue, which requires a lot of synchronization between the tasks. Since we are going to repeat the algorithm several times, we can use instead a simple and efficient sequential Dijkstra implementation and distribute the source vertices over the available processing elements (PEs).

Data distribution

The data involved in this computation are the (static) description of the graph and the (dynamic) current shortest path spanning tree built by Dijkstra's algorithm. The latter is PE-private. The former, being pure *read only* data could be distributed over the PEs without having to worry about data-consistency issues. However, the way this data is accessed by the single source shortest paths method generates a tremendous amount of small communications that can bring even the most efficient communication system to its knees. The amount of memory that comes with today's machines allows in this case to simply duplicate the graph description on every PE, ensuring only local memory access.

There remains only a little cooperation between the PEs to exchange information about which nodes have already been computed, what are the current best values of the bounds, and whether or not to stop the computation. Various approaches have been used to address these issues.

On the SGI Origin 2000, the shared memory multithreading programming model was used first. It is a very fast and convenient way to turn a sequential code into a parallel code one step at a time. The results with a shared adjacency list were not satisfying, but once it was duplicated on all the threads, quasi-linear speedup could be achieved. Table 2 shows the wall clock times in seconds on the Origin 2000 for a sparse graph of 2423 vertices and 14538 edges with random weights.

Number of PEs	1	2	4	8	16
Time (seconds)	433	223	113	58	30
Speedup (T_n/T_1)	1.00	1.94	3.83	7.46	14.43
Efficiency (T_n/nT_1)	1.00	0.97	0.96	0.95	0.90

Table 2: Multithreading on the Origin 2000

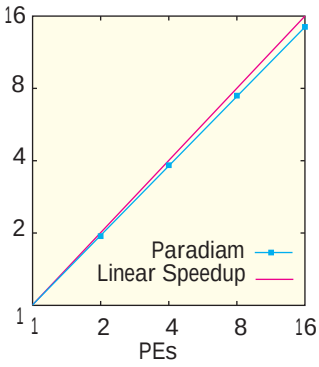

Fig. 2 – Speedup

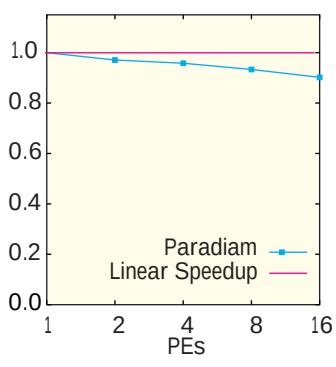

Fig. 3 – Efficiency

Then we ported the code to MPI, with two different approaches. The first one simulates the multithreading code, substituting `Allreduce` operations for the global variables. This gave good results on the Origin 2000 and on the Swiss-T0, where the tasks are usually well synchronized and the `Allreduce` operation is supported by efficient hardware. On the Origin 2000, the results matched those of the multithreading code. This is a hint that communication is not the limiting factor in this application. Table 3 gives the wall clock times in seconds for the polytopal graphs. The Origin 2000 is better on medium problems where all the data fit in cache, while the Swiss-T0 has a better main memory bandwidth for larger problems.

d		10	11	12	13
1 PE	Origin	40	**210**	**2052**	–
	T0	**29**	285	3300	–
2 PEs	Origin	20	**105**	**1037**	11354
	T0	**18**	152	1687	**9514**
4 PEs	Origin	12	**54**	**664**	6333
	T0	**10**	79	859	**5122**
8 PEs	Origin	**6**	**32**	**382**	3209
	T0	8	42	441	**2504**
	O2	19	99	1080	6153

Table 3 – MPI on the Origin 2000 the Swiss-T0 and the O2

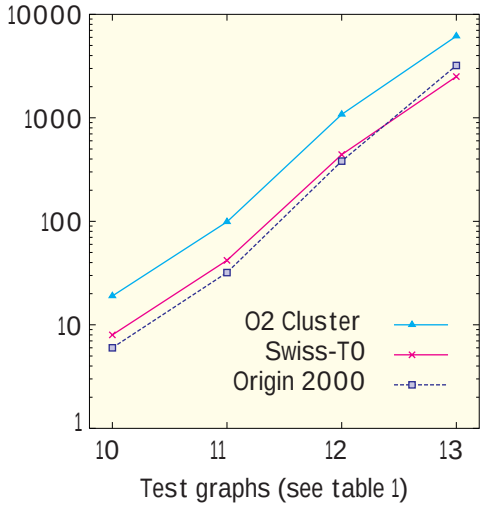

Fig. 4 – Running times on 8 PEs

The same code was then run on a cluster of 20 SGI O2 workstations. Too much time was lost in synchronizing the tasks and performing the `Allreduce` operation on the slow 10Mbit/s Ethernet. To achieve better performance in this case, a master-slave model was developed where one task is responsible for allocating the sources to the slaves, gathering the results and updating the bounds. The buffered, non-blocking communications offered by MPI allow in this case keeping the slaves busy 100% of the time, even in a heterogeneous environment with a poor interconnection network. The downside of this approach is of course the need for an extra computing node for the master task, but it is so light (both CPU and memory) that any desktop PC with a decent network connection to the cluster is a good candidate to host it. It was also possible to map the master task and one slave on one physical PE without a significant performance drop.

As we expected, the implementation on the cluster of O2:

- scales well only on large problems;
- suffers from very high startup time on small problems;
- is generally less efficient than the Origin 2000 or the Swiss-T0, because of the respective single-PE performances;
- is way behind in non-measurable terms of *ease of use*.

Nevertheless, as all the teams worldwide who are putting together piles of cheap Linux boxes will assert, it proved that clusters remain a good alternative to the *big irons* for applications that do *not* require efficient communications.

CONCLUSIONS

Even if the need for an efficient way to compute graph diameters is real, this remains a relatively light application. Careful study of the available methods allows selecting those that will give the best results in a parallel context. In this case, it was possible to design the code in a way that reduces communications to a strict minimum. Therefore, all the conclusions that can be drawn from this study do not pay justice to the efficiency (or lack thereof) of the intercommunication networks of the parallel machines or clusters.

Furthermore, the pattern of the memory accesses in a typical run of Dijkstra's algorithm has to be taken into account when comparing the performances. In all cases, the on-chip caches (in the order of 10 to 100 Kb) are defeated, so the high clock frequencies have only a limited influence on the overall performances. The large 2nd level cache of the Origin 2000 PEs give this machine a significant advantage on medium size problems, while the nodes of the Swiss-T0 seem to have a better main memory bandwidth for larger graphs [5]. Talking about the Swiss-T0 prototype, its high bandwidth EasyNet bus was not used at all here, partly because the required parts of MPI were not yet available, and partly because the very limited needs in communications did not justify it. But other applications will certainly benefit from this technology !

We would like to thank Ralf Gruber, Director of CAPA, for his valuable advice and for granting us access to the Swiss-T0 prototype, and Jean-Claude Berney, DDI of the DMA, for providing us exclusive access to the SGI O2 of the student's lab at the DMA.

REFERENCES

[1] G. S. Brodal, J. L. Träff, and C. D. Zaroliagis. *A parallel priority data structure with applications.* In Proc. 11th International Parallel Processing Symposium, pages 689-693, 1997.

[2] G. B. Dantzig. *Linear Programming and Extensions.* Princeton University Press, 1963.

[3] J.-A. Ferrez, K. Fukuda, and Th. M. Liebling. *Parallel Computation of the Diameter of a Graph*. to appear in High Performance Computing Systems and Applications. proc. 12th Annual International Symposium on High Performance Computing Systems and Applications (HPCS'98), Kluwer, 1998.

[4] J.-A. Ferrez. Web page with links and codes about the graph diameter computation. ***http://rosowww.epfl.ch/jaf/pages/paradiam.html***.

[5] J.-A. Ferrez. *A quick evaluation of the machines available at DMA with HINT*. ***http://rosowww.epfl.ch/jaf/HINT/***.

[6] A. V. Goldberg and R. E. Tarjan. *Expected performance of Dijkstra's shortest path algorithm*. Technical Report 96-070, NEC Research Institute, Inc., 1996.

[7] P. McMullen and G. C. Shephard. *Convex polytopes and the upper bound conjecture*. Cambridge University Press, 1971.

[8] R. Seidel. *On the all-pairs-shortest-path problem*. In Proc. 24th Annu. ACM Sympos. Theory Comput., pages 745-749, 1992. ■

UN SYSTÈME D'INFORMATION POUR LA GESTION INTÉGRÉE DES COÛTS DANS LA CONSTRUCTION

ABDELILAH ZERTITI, EPFL - DÉPARTEMENT DE GÉNIE CIVIL, LEM

Information systems are actually a strategic element for the enterprise. They play a crucial role for organizational efficiency, as well as for competitiveness. In the present state of the construction industry, it becomes a vital component to enable the deployment of a rational management system. This management has to give answers to two principals problems: on one hand the increased complexity of the construction projects, and on the other hand, the need of a closer control of the critical project's parameters (Quality, Cost, Time). This implies control of detailed, dense and highly integrated information, which, in turn, implies new concepts and a large computer resources. Because of the complexity of this problem, the study, presented in this paper, is focused on the description of a tool designed for an integrated cost management in the construction industry: Ges2000[1].

Les systèmes d'information s'imposent actuellement comme un élément stratégique pour l'entreprise tant dans sa recherche d'efficacité organisationnelle que dans sa recherche de compétitivité sur les marchés. Dans le contexte difficile actuel du secteur de la construction, ils deviennent un point d'appui incontournable pour accompagner la mise en place d'une gestion rationnelle. Cette gestion doit répondre, d'une part, à la complexité croissante des projets de construction et, d'autre part, au besoin d'un contrôle plus fin des paramètres critiques d'un projet (qualité, coût, délai). Cela implique la maîtrise d'une information très détaillée, dense et hautement intégrée, nécessitant de nouvelles approches de conception et de grandes ressources d'informatique de gestion. Le problème étant vaste et complexe, et comme premier pas vers une intégration des autres aspects de gestion, notre étude s'est concentrée sur la conception et le développement d'un outil pour la gestion intégrée des coûts dans la construction : Ges2000[1] .

1 Ges2000 a été développé sous la direction du professeur F.-L. Perret, en étroite collaboration avec le bureau d'ingénieurs-conseils CETP. La coordination du projet a été assurée par Dr Ph. Wieser. Le développement de la partie soumission a été réalisé respectivement par MM. A. Essade et J. Laventure. Le modèle conceptuel global Entité-Association ainsi que la conception et le développement des autres modules (calcul des frais par éléments, devis, contrat et suivi de la réalisation et des paiements) ont été réalisés dans le cadre de la thèse présentée par A. Zertiti sous le titre "Qualité Totale et réseaux stratégiques dans la construction, EPFL 1997".

GESTION INTÉGRÉE DES COÛTS DANS LA CONSTRUCTION

L'environnement d'une opération de construction est caractérisé par la multiplicité de partenaires et la diversité des tâches à accomplir, couvrant l'ensemble des activités liées à la conception, à la réalisation et à l'exploitation d'un ouvrage. Comme dans tout projet, elles sont de natures diverses et font appel à des données tout d'abord floues et incertaines qui, progressivement, vont se préciser en fonction de l'âge du projet et de son degré de maturation.

Si traditionnellement ces tâches faisaient l'objet de responsabilités séparées et bien délimitées, pour lesquelles des mandats précis étaient attribués, les tendances actuelles, confirmées par plusieurs demandes explicites de grands maîtres d'ouvrage, vont vers une intégration totale de l'ensemble de ces tâches.

Il en résulte, de manière symétrique, une nécessité accrue d'améliorer la qualité, la disponibilité et la pérennité des données à travers des systèmes d'information intégrés permettant de refléter le plus fidèlement possible cette tendance à l'intégration des tâches. Cette intégration permettra, du fait de la grande cohésion des données, une confrontation aisée et permanente de celles-ci, des possibilités accrues d'exploitation statistique des coûts et une ventilation adaptée des frais selon les besoins.

DESCRIPTION DU SYSTÈME D'INFORMATION

LIBRAIRIES DE RÉFÉRENCE

Il s'agit de l'ensemble des catalogues standard codifiant les prestations sur le marché de la construction suisse. Cette codification est très détaillée; l'ensemble des catalogues représente une masse impressionnante d'informations, qui ne cesse de croître. En plus, ces standards ne sont pas toujours utilisés et l'on continue à faire appels aux catalogues propres. Il a fallu, dès lors, proposer une gestion homogène des standards suisses en cohabitation avec d'autres catalogues et d'autres standards à travers un concept simple et facile à mettre en œuvre.

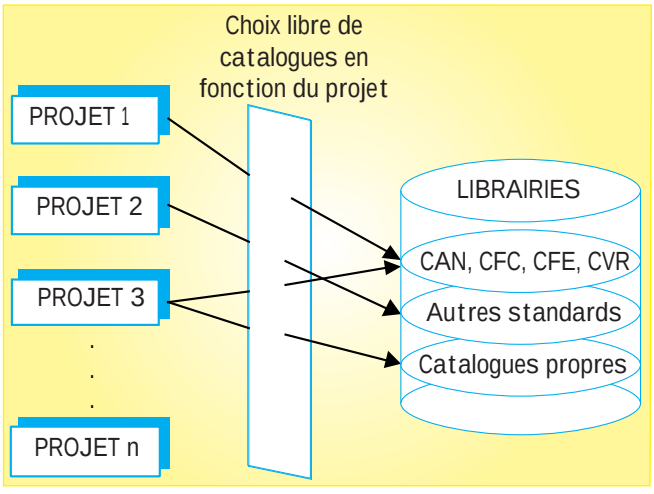

Fig. 1 – Gestion des librairies de catalogues

La notion de librairies de référence est une approche qui consiste à classer sous forme de librairies tous les catalogues de référence standard ou propres. A chaque mandat, une librairie est associée comme un attribut propre du projet. Cela permet, selon les méthodes du travail du projeteur et les exigences de la maîtrise d'ouvrage, un choix souple et adapté à chaque projet. Ainsi, il devient possible de mener en parallèle des projets utilisant des standards différents.

DÉCOUPAGE DU PROJET

La conception du système a privilégié son adaptation à des problèmes multiples et variés dépendant des conditions propres de chaque projet, ainsi qu'une souplesse d'utilisation en matière de saisie, de visualisation, d'agrégation ou de désagrégation de l'information selon tous critères de recherche. Chaque prestation peut être définie de manière précise et détaillée (fig. 2). En l'associant à un contrat (lot ou entreprise), cette structure permet un contrôle rigoureux des coûts, à travers la définition de divers types de structures utiles pour ventiler l'information. Les dépenses peuvent être ventilées, entre autres: par contrats (lots), par ouvrage, par entreprise pouvant cumuler plusieurs contrats, par groupes comptables, etc.

En plus de ces différents découpages que nécessite le contrôle des prestations, un découpage spatial est nécessaire. La finesse de celui-ci dépendra de la taille du projet et des besoins du suivi. L'outil intègre la possibilité de subdivision de l'ouvrage par objet et par localisation. La subdivision par objet permet de décomposer en parties d'ouvrage et d'attribuer à celles-ci les différentes prestations y relatives. La subdivision par localisation permet, comme son nom l'indique, de déterminer à l'intérieur d'un projet l'emplacement des prestations.

A noter, que pour tenir compte de la diversité des langues suisses et de l'activité économique de plus en plus globale et internationale, la langue a été conçue comme un attribut paramétrable du projet, permettant ainsi une description des prestations adaptée à l'environnement linguistique du projet.

LES PRÉVISIONS (DEVIS)

Deux approches d'évaluation sont intégrées dans l'outil. La première, classique, est axée sur l'exécution des travaux, tandis que la deuxième, récente, est plutôt conceptuelle.

L'approche axée sur l'exécution des travaux part d'une décomposition des postes de frais en fonction des prestations à exécuter, adoptant pour cela une systématique d'évaluation Top-Down. Plusieurs variantes peuvent être étudiées et proposées avant un choix définitif.

L'approche conceptuelle, se situe à l'amont de la première. Elle part d'abord d'une estimation des coûts par éléments fonctionnels de l'ouvrage, en faisant appel à des catalogues de prix de référence. Sur la base de ces valeurs de référence, des estimations peuvent être rapidement établies, depuis des stades sommaires jusqu'aux niveaux les plus détaillés. On retrouve la même systématique Top-Down ci-dessus.

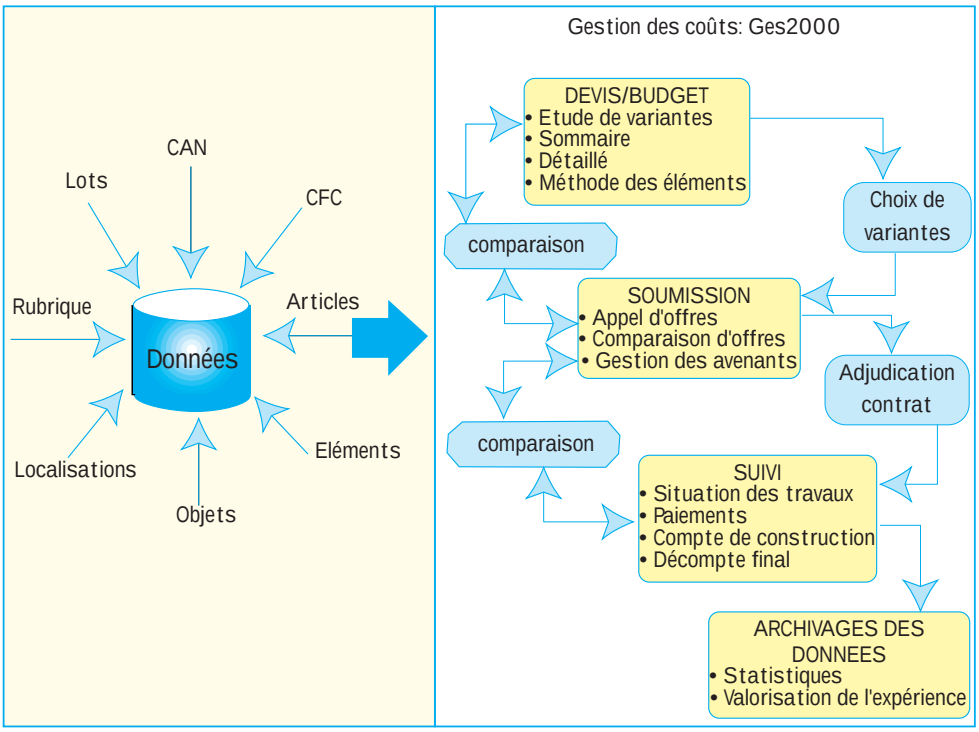

Fig. 2 – Structure des données

Le passage d'une logique conceptuelle à une logique d'exécution pour les appels d'offre se fait à travers une matrice associant les éléments et les différentes catégories de travaux nécessaires à leur réalisation.

Le temps est un élément important dans les prévisions financières. Dans les projets de construction, le temps séparant les études de la réalisation proprement dite peut être assez long, et justifier, parfois, le recours à un mécanisme d'indexation pour l'actualisation des montants prévus. L'outil propose une gestion souple et riche en intégrant un choix varié d'indices standard ou propres. Cette souplesse rend possible la combinaison, selon les besoins et les travaux concernés, de plusieurs indices différents au sein d'un même projet. Ceci, avec une mise à jour automatique, suivant la fréquence de la publication des nouveaux indices.

Gestion des contrats et des comptes spéciaux

En plus des soumissions et des comparaisons d'offres, c'est surtout la gestion des contrats qui présente une originalité. L'adjudication d'une soumission à une entreprise donne lieu à un contrat. Différentes déductions et retenues sont appliquées à ce contrat. Elles peuvent être très importantes et être affectées à la couverture de dépenses les plus diverses. Autrement dit, elles peuvent être utilisées pour effectuer des paiements, des transferts et y prélever d'autres déductions. Cela pose de très sérieux problèmes de gestion au niveau de la souplesse d'utilisation et de transparence.

Pour pouvoir assurer une gestion homogène des différents flux financiers et une transparence optimale des origines et destinations des mouvements, nous avons choisi une modélisation à travers la notion globale de compte. Autrement dit, et au même titre qu'un contrat, toutes les déductions (pro rata, rabais, garantie, etc.) sont gérées pour elles-mêmes et peuvent être considérées comme des comptes sur lesquels il devient possible d'effectuer directement des prélèvements et des paiements.

Au niveau de la définition même des contrats, l'outil permet de tenir compte des types de contrats les plus divers : selon des arrangements forfaitaires, selon un plan de paiement avec facture finale ou selon des attachements et des situations périodiques détaillées. En plus, le contrat peut intégrer tous les types de déductions et conditions répondant aux environnements contractuels les plus variables. On a, par exemple, la possibilité d'utiliser un rabais de régie échelonné, des déductions en garantie par palier, un calcul automatique des hausses selon le contrat (formule mathématique ou montant exogène), l'application d'un escompte, des conditions de primes et de pénalités, etc.

Le suivi financier

Il est important pour les paiements et le contrôle des coûts de définir dès le début du projet et avec grand soin les divers types de structures utiles pour ventiler l'information concernant les dépenses. Cela permet un suivi fin et une imputation de ces dépenses sur des centres de frais de la manière la plus transparente possible. C'est pourquoi nous avons tenu à faire la distinction entre quatre types de factures à savoir:

Métré: tout ce qui se rapporte à une soumission adjugée avec ses avenants.
Régie: tout ce qui se rapporte à des travaux facturés en fonction des heures et des matériaux utilisés.
Hausse: tout ce qui est conséquence d'un renchérissement du coût de la vie dans le domaine de la construction.
Divers: tout ce qui n'appartient pas aux trois rubriques précédentes.

Tous les relevés périodiques de travaux et de dépenses reflètent cette structure. Ceci est réalisé à travers l'association systématique d'une rubrique, permettant cette différenciation en relation avec la structure du devis. Une *situation* permet de gérer les relations financières de chaque compte et peut concerner, comme cela a été dit, une soumission adjugée ou un partenaire en dehors d'une adjudication.

A chaque *situation* correspond un relevé détaillé des quantités réalisées au niveau des prestations prévues dans le

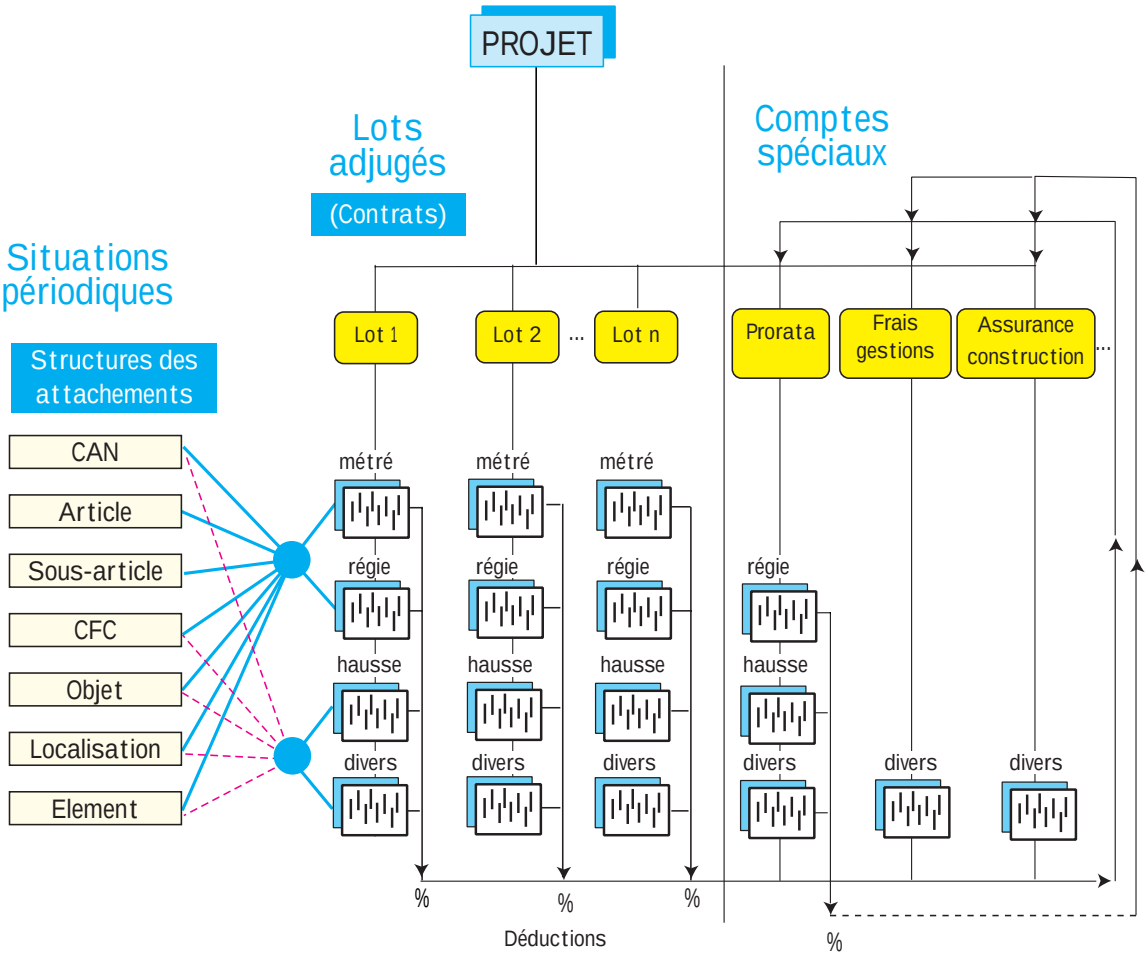

Fig. 3 – Structure des dépenses et organisation des comptes

contrat sur la base de prix unitaires ou sur la base d'autres tarifs horaires prévus au niveau de la régie.

Pour les autres dépenses, les différents relevés correspondant à de petits travaux ou fournitures ou à des hausses survenues en cours de travaux peuvent être, selon le degré de finesse voulu, ventilés par catégorie de travaux, par poste comptable, par élément, par objet par localisation ou une combinaison de l'ensemble.

Sur la base des relevés périodiques des *situations* ou sur la base des acomptes forfaitaires prévus, un bon de paiement est édité. Le calcul des différents montants se fait selon les conditions contractuelles prévues. Chaque paiement effectué est sauvegardé avec toutes les conditions contractuelles qui s'y appliquent. Ce qui permet, malgré l'évolution des contrats et les modifications de certaines conditions, de retrouver à chaque instant l'environnement contractuel d'un paiement donné.

Pendant la phase de réalisation, la densité de l'information sur la chaîne de valeur est la plus grande. Les flux de paiements et des documents y relatifs sont importants exigeant un contrôle rigoureux des *situations* et des paiements. Ceci est réalisé à travers la définition de processus d'acquisition et de validation de l'information par les responsables et un système de privilèges hiérarchiques.

CONCLUSIONS

Les systèmes d'information s'imposent actuellement comme un élément stratégique pour les entreprises. Dans le domaine de la construction, la tendance actuelle vers une intégration totale de l'ensemble des tâches exige, de manière symétrique, la conception de systèmes d'information intégrés, visant à refléter le plus fidèlement possible cette tendance. Cela implique la maîtrise d'une information très détaillée, dense et hautement intégrée, nécessitant une conception adaptée et de grandes ressources d'informatiques de gestion. Nous avons présenté, dans cet article, une conception de ce que l'on pourrait considérer comme une première étape sur la voie de l'intégration de l'information sur la chaîne de valeur d'un projet de construction. Elle concerne principalement la gestion des coûts depuis les premières ébauches jusqu'au décompte final. L'article a décrit la structure et la nature des données à intégrer pour une gestion rationnelle et une exploitation optimum de l'information. ■

A HIERARCHICAL COMPUTER SIMULATION MODEL FOR THE EVOLUTION OF THE MICROSTRUCTURE PRODUCED BY DISPLACEMENT CASCADES IN METALS

Abderrahim Almazouzi & Maximo Victoria, EPFL-CRPP-EURATOM Association, Fusion Technology Materials and Maria José Caturla & Tomas Diaz de la Rubia, Lawrence Livermore National Laboratory, Chemistry and Materials Science Division

La prédiction de l'évolution de la microstructure des matériaux, lorsque ceux-ci sont irradiés avec des particules capables de créer des cascades, nécessite une profonde évaluation des caractéristiques de la production des défauts dans les cascades aussi bien que des propriétés des défauts ponctuels et de leurs amas. Expérimentalement, il existe des problèmes techniques insurmontables étant donné l'échelle microscopique de ces défauts (<1nm) d'une part et la très courte durée de vie des cascades (<10ps) d'autre part. Une des quantités clefs nécessaires, permettant d'émettre des prédictions quantitatives et de pouvoir comparer les résultats de simulation avec les résultats expérimentaux, est la fraction des défauts libres mobiles. Cette fraction dépend fortement de l'état initial des dégâts déposés par irradiation, de la mobilité relative des différents défauts produits ainsi que de leurs interactions avec la microstructure initiale du matériau irradié. Différents modèles théoriques basés principalement sur l'idée de la non-équipartition des défauts produits lors de la cascade (Production Bias Model) ont été appliqués avec succès pour expliquer le comportement des métaux purs irradiés sous différentes conditions. Cependant ces modèles négligent encore un ingrédient de première importance qui est celui relatif aux mécanismes de diffusion des défauts produits par irradiation.

Récemment, un modèle hiérarchique du simulation par ordinateur prenant en compte les propriétés locales et instantanées des défauts a été développé. Ce modèle comprenant deux méthodes de simulation numérique, à savoir la dynamique moléculaire et Monte Carlo, est utilisé pour étudier l'évolution des dégâts d'irradiation dans un métal pur Al.

The prediction of the microstructure evolution of metals under cascade forming irradiation, necessitates an in-depth evaluation of the features of the defect production in the cascades and the properties of point defects and clusters. Experimentally, there are insurmountable technical problems to investigate these properties because of the microscopic scale of these defects (<1nm) and the shortness of the cascade life time (< 10ps). One of the key quantities required to make quantitative predictions and to compare simulation results to those of experiments is the fraction of freely migrating defects (FMD). This fraction is strongly dependent on the form of primary damage state, the relative mobility of the various clusters produced, and their interaction with the initial microstructure. Although analytical models based on the production bias model have been applied successfully to explain the behaviour of pure metals under different irradiation conditions, they still lack an important ingredient, namely that of the diffusion mechanism of the defects produced.

A hierarchical model of computer simulation which takes into account the local and temporal features of the defects has been developed. This model which is based on a combination of two methods namely Molecular Dynamics and Monte Carlo has been employed to study the damage evolution in pure Al.

Introduction

The main interest of the Fusion Technology Materials(FTM) group at EPFL-CRPP is to investigate the effects of the damage produced by radiation in metals and alloys, especially the candidate materials for structural components of the future fusion reactor. This group, located at Paul Scherrer Institute in Villigen uses the high energy proton beam (590MeV) in the **P**roton **Ir**radiation **Ex**periment (PIREX) in order to simulate experimentally the effects of the 14MeV neutrons that are the products of the fusion reaction between D and T.

The energetic particles cause a number of effects on the structural materials of the reactor, amongst them:

- Radiation hardening and eventual degradation of the mechanical properties (loss of ductility and fracture toughness).
- The generation of radioactive impurities and gases (H and He) inside the exposed materials by nuclear reactions can eventually induce the formation of gas bubbles.
- The creation and growth of voids due to the accumulation of point defects will lead to swelling.

These effects are the main factors limiting the choice of candidate materials. The residual radioactivity of a large amount of exposed material is also a concern and will govern the handling methods, dictate the storage periods

and the overall waste management and recycling scenarios.

The research activity of the group is concentrated mainly in the development and testing of new materials where some or all of the deleterious above cited effects can be hindered or at least minimised. As an underlying part of this development, fundamental aspects of the irradiation damage in solids are an essential part of the overall research program.

It is within this frame that computer simulation activities in this field were first established in 1988. Three main areas were initially covered:

(a) semi-empirical quantum chemistry [1-5]
(b) ultra-fast processes in solids [6-10]
(c) radiation damage in intermetallic compounds and metals [11-15].

The irradiation of metals with cascade forming particles produces in addition to Frenkel pairs (simple vacancy-interstitial pairs) defect clusters, which may form either during the evolution of a single cascade, or during the accumulation of the damage, i.e. during the in-service of the materials in an irradiation environment. The actual defect concentration and its correlation within a cascade strongly influences the formation of large defect clusters during annealing after low temperature irradiation or during irradiation at elevated temperatures. Such defect clusters are of great technological relevance as they are stable up to higher temperatures and can change drastically the mechanical properties.

Direct observation of this irradiation induced defect distribution is difficult and scarce due to the very small size of the clusters on one hand and to their very rapid escape from their nascent cascade on the other. The first stages of the cascade development including the binary collision phase as well as the thermal spike have been successfully described by means of Molecular Dynamics (MD) using semi-empirical potentials. In fact, within last decade the correlation between the Primary Knock Atom (PKA) energy and type and amount of defects as a function of the material and its features (mass, density, temperature and structure) has been largely explored [16]. The overall picture is that the cascade duration is in the tenth of picoseconds range, the zone heated up is the nm scale when considering the isolated cascade and its structure consisting of a vacancy rich region (depleted zone) surrounded by a cloud of interstitials. It has also been shown that these defects may agglomerate to form very stable clusters within the thermal spike duration. This picture has been confirmed in some cases by experimental observations. Although such simulations are very important to assess the primary damage stage, they are very limited in the time (ps) and space scales (tens of nm) because of the large computing processing media needed.

One of the key quantities required to make quantitative predictions of microstructure evolution of irradiated materials and to compare simulation results to those of experiments is the fraction of defects that escapes the nascent cascade region, the so-called the freely migrating defects (FMD). This fraction is strongly dependent on the form of primary damage state, the relative mobility of the various clusters produced and their interaction with the initial microstructure [17]. In order to describe the long term evolution of the defects produced by the cascade, analytical models based on the production bias model have been applied successfully to explain the behaviour of pure fcc metals under different irradiation conditions [18]. However, they still lack an important ingredient, namely that of the diffusion mechanism of the defect distribution produced after the quench of high energy cascade. Recently, a hierarchical computer simulation model including Molecular dynamics (MD) and Kinetic Monte Carlo (KMC) methods has been successfully employed to extract the FMD fraction in bcc Fe [19] and fcc Au [20]. The same approach has also been applied to predict the ion implantation behaviour in Si [21].

In the frame of a collaboration program with the Lawrence Livermore National Laboratory (USA), this new hierachical computational model has been used to describe the long term evolution of the defects produced by the cascade. It is based mainly on the powerfulness of the Molecular Dynamics (MD) method when using empirical interatomic potentials and a vectorized and parallelised computer code, and in the flexibility of Kinetic Monte Carlo (KMC). The two methods are used in the spirit of bridging the space and time scale proper to each of them. While the MD deals exclusively with interactions at the atomic distances and over a time scale of few tenth of picosecond, the KMC allows a full atomistic study over extended time scale (up to hours) and space scale up to micrometer. A necessary condition is the use of a self consistent data set for the energetics of defects and their diffusion, which may be obtained from MD.

The goal of such type of modelling is the assessment of the lifetime of materials when they are used in an irradiation environment. (see fig. 1 for the evolution of radiation damage).

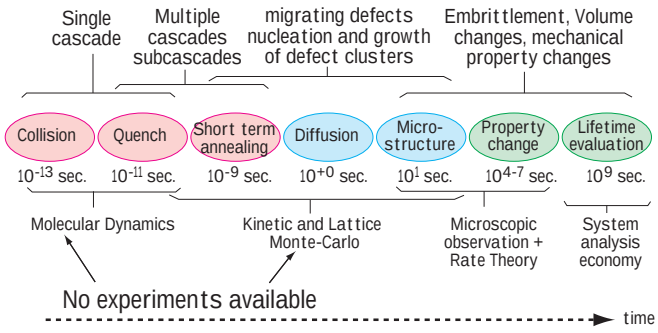

Fig. 1 – Radiation Damage evolution

In fact, irradiation of solids with either heavy ions or high energy particles is a subject of interest from the technological as well as the fundamental point of view: ion implantation followed by high temperature annealing of the damage is a standard processing step in manufacturing silicon VLSI devices [22]. Even at the low doses used in some manufacturing steps, implantation introduces disorder and defects into the lattice. Another example is that of

nuclear power stations, where irradiation-assisted stress corrosion cracking (IASCC) and the influence of irradiation on fatigue, mechanical wear, erosion, embrittlement, creep, and stress relaxation have been recognised as the primary ageing-related degradation mechanisms [23].

This present report describes the MD simulation of the self ion irradiation in Al for PKA's up to 15 keV. The MD method was also used to extract the energetics (migration, formation and binding energies) of the defects. The long term simulation of the damage evolution was performed using a Kinetic Monte Carlo (KMC) model. The aim of the research program is to elucidate the atomistic mechanisms that occur during the microstructural evolution under irradiation over observable length and time scales as well as in a wide range of temperatures.

Simulation Methodology

Molecular Dynamics Simulations

Simulations of low energy recoil events, high energy displacement cascades, defect energetics (formation and binding) and kinetics (migration) were all carried out with the MD simulation code MDCASK [24]. The MDCASK code was developed to run in a CRAY T3D system by the National Energy Research Supercomputer Center at Lawrence Livermore National Laboratory and implemented in the T3D of the Ecole Polytechnique Fédérale de Lausanne with the help of T3D users support team in the Service Informatique Central (SIC). The domain decomposition used in the program is based on the link cell method, which consists in the partition of the cell into smaller cells, each of them with dimensions slightly bigger than the cut off of the interatomic potential. With this distribution, the possible neighbours of each atom are calculated only from these atoms that are in its same link cell or on the 26 other link cells surrounding it .

The calculation of forces and energies can be done then independently in each processor. However, some information needs to be shared between processors, specially in the calculation of forces and energies for atoms close to the boundary between two processors. The communication between processors is done using the message passing library PVM developed in Oak Ridge National Laboratory. The performance of this code is about $2.\ 10^{-5}$ sec/step/atom in a 64 PE's node. The classical equations of motion are integrated numerically using a fourth order predictor-corrector scheme for each particle in the system. The time step is automatically adjusted so that the criteria for a maximum displacement of $0.1a_0$ (where a_0 is the lattice distance) and a maximum allowed fractional energy change of 1 eV/at. are satisfied at any time step for each atom in the crystal. Almost all the calculations of the present study have been performed using the constant volume and periodic boundary conditions. The motion of the atoms in the link cells adjacent to the cell boundary is governed by the Langevin equation, in order to control the temperature of the crystal.

The Al interatomic potential used is the one developed by Ercolessi and Adams [25] and which is based on the fitting to ab-initio calculated atomic forces of many atomic configurations at finite temperature. The calculation of many of the properties of Al with this potential has been shown to reproduce the experimental values especially the stacking fault energy, which is much higher than that of any other fcc metal. This potential has been modified following the procedure described by Calder and Bacon [29] in order to properly describe the short range high energy scattering.

Kinetic Monte Carlo model

The model used in this investigation (BIGMAG code) is inspired from the one developed recently by Heinish [26] to perform stochastic annealing of the defects produced by a 25 keV cascade in Cu. In the BIGMAC code [27], the single point defects (interstitials, vacancies), and the extended defects (vacancy and interstitial clusters) are considered as individual objects with a basic set of properties such as size, jump rates and dissociation rates. There is no underlying lattice in the Monte Carlo box, since the defects are considered to jump randomly in one or three dimensions. Thus, a relatively big volume, up to cubic micrometer can be simulated in a modern workstation (e.g. SGI-R10000) within reasonable CPU time. The values of the different reaction rates, such as those for the formation or dissociation of clusters, as well as the jumping rates of different species are given by the binding and migration energies, respectively. The jump distance for any type of defect present is set equal to the first nearest neighbour (0.285 nm) for the simulation in Al. The initial stage of the simulation consists of the three-dimensional distribution of the defects that are produced by the MD simulation of the displacement cascade. Events are picked-up according to the reaction rates and the number of objects of any one type. The object selected is either placed in a new location (in case of migration or dissociation from a cluster) or changed into another type of object (in case of clustering) or deleted from the system (in case it is trapped by a sink in its migration path). The former and new neighbors of the chosen object are modified accordingly. This last step is the most CPU costly in the simulation. The time is incremented after each event considering the total reaction rates. The real time is therefore linked to a KMC simulation step through the relation:

$$t = \frac{\delta^2}{6D} \qquad (1)$$

where: t, δ and D are the real time, the jump distance and the diffusivity, respectively. The code uses temperature dependent diffusivities to calculate jump probabilities and considers the migration of point defects and clusters distinctly.

The data required to carry out the KMC simulations are the self diffusivities of defects (single and clusters of different sizes), the spontaneous recombination volume of vacancies and interstitials and the binding energy and capture radius of point defects and clusters, in addition to the jump distance. The temperature dependence of the defect diffusivity is described by the Arrhenius law as:

$$D = D_0 \exp\left(-\frac{E_m}{kT}\right) \quad (2)$$

where D is the defect diffusion coefficient, D_0 is the pre-exponential (or frequency) factor, E_m, k and T are the migration energy, the Boltzmann constant and the temperature, respectively. In the present investigation, the values of E_m, D_0 as well as the binding energies for clusters of different type and size are obtained from MD simulations of diffusion at finite temperature and defect relaxation methods, respectively.

RESULTS AND DISCUSSION

CASCADE SIMULATIONS

The displacement cascades were simulated for PKA energies ranging from 100 eV up to 15 keV. The size of the Al simulation matrix, which was relaxed at 10 K for about 1ps before launching the PKA, has been varied between 186424 up to 500000 atoms depending on the PKA energy that has to be simulated. Care was taken in the choice of the orientation and the position of the PKA, in order to make sure that all displacements happen mainly in the matrix centre. The temperature of the first boundary layers was set at 10 K and controlled with the Langevin equation during the whole process until the temperature of the entire box matrix reached 10 K.

Both kinetic and potential energy as well as the mean square displacement of all atoms constituting the matrix have been monitored. The configuration data files have been analysed in terms of the distribution of the defects produced.

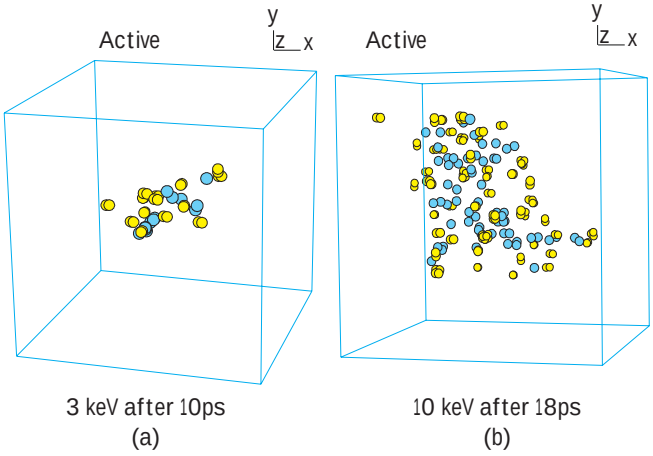

Fig. 2 – Defect distribution produced by self ions of (a) 3 and (b) 10 keV in Al at 10 K. Blue circles represent vacancies and the yellow ones interstitials

Fig 2. shows the three-dimensional distribution of the vacancies (blue spheres) and interstitials (yellow spheres) obtained at the end of the quenching process of the thermal spikes produced by an Al PKA of 3 keV at 10 ps (Fig2a) and 10 keV at 18 ps (Fig2b). In the 3 keV case the above described displacement cascade picture, e.g. a depleted zone surrounded by an interstitial cloud, still holds. In contrast, for the 10 keV cascade an almost homogeneously disperse distribution of defects is observed, i.e. there are no distinguishable zones.

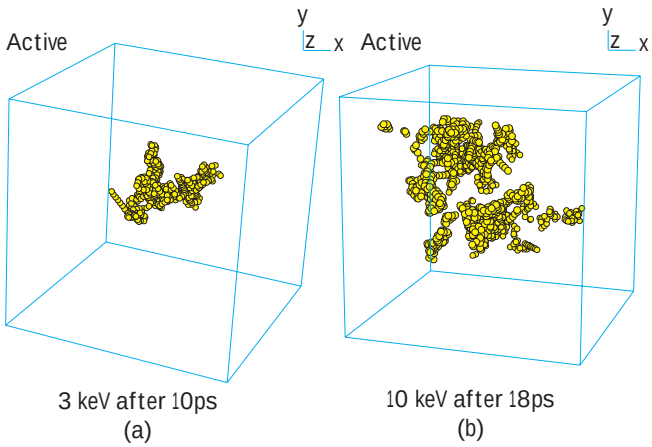

Fig. 3 – Displacement cascade produced by self ions of (a) 3 and (b) 10 keV in Al at 10 K

Fig.3 shows the 3D distribution of all atoms displaced during the evolution of both cascades (3 keV and 10 keV) and gives an insight into the difference between the defect distribution, at different energies. The low energy PKA induced only a single cascade, while the higher energy one produced at least two distinguishable zones with a high energy density. It can be stated therefore, that the earlier picture of cascade development holds only in the case the PKA energy is not enough to produce subcascades, i.e. below the threshold of subcascade formation.

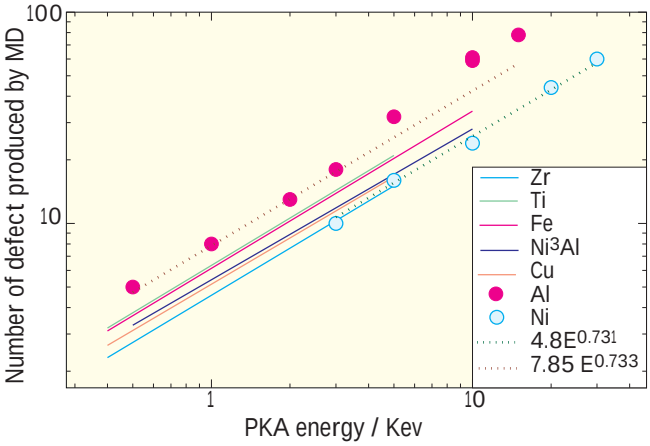

Fig.4 – Number of Frenkel pairs calculated with MD in Al at 10 K compared to values in other materials at 100 K, ref. [16]

The dependence of the number of the defects produced on the PKA energy is plotted in fig.4, together with our recent results in Ni [28] and those reported by Bacon et al [29] for other metals. It is noticeable that the amount of the defects produced in Al is higher than in the other metals, due most probably to its low density.

Furthermore, when using the power law proposed by Bacon et al [29] ($N=AE^m$ where N is the number of defects produced by a PKA of energy E, A and m are fitting parameters) two regions are visible: the first one for energies lower that 5 keV, which can be described by a power law with fitting parameters comparable to those for the other metals (A=7.85 and m=0.733). The second region is above about 5 keV and shows a clear upward deviation from the power law fit. This difference can be explained by taking into account that the energy is over the threshold of subcascade formation. This result is comparable to that reported for simulations of 40 keV PKA cascades in Fe by Stoller et al. [30].

An equivalent result has also been obtained from simulations using the liquid droplet model using the MARLOWE binary collision code in Al [31].

As for the defect production efficiency in Al, fig. 5 shows the normalised defect production rate (ratio of the number of MD calculated defects to the number of defects calculated using the NRT formula [32]). This ratio saturates at about 20 to 25 % for energies higher than 2 keV in agreement with the values obtained in all the simulations performed by different groups so far in different metallic materials [33]. In the same fig. 5, four additional curves are plotted to demonstrate the intracascade clustering efficiency in Al. It is worth noticing at this point, that more that 90 % of the defects produced at all the PKA energies investigated are singles (isolated defects). The rest is agglomerated in several different small clusters (see fig. 2) most of which are of interstitial type. The biggest cluster found in this work contains 5 interstitial defects.

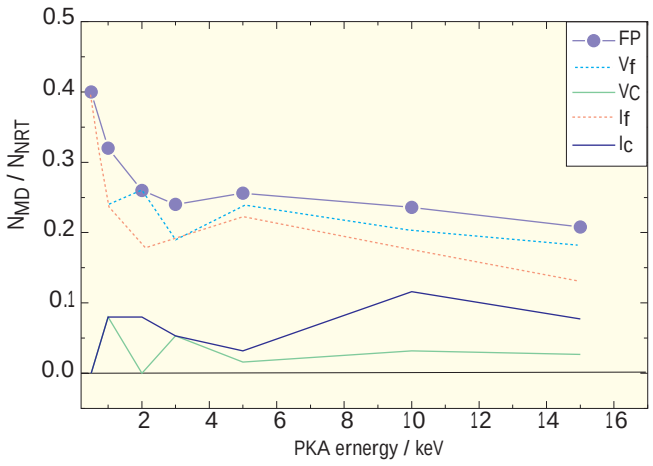

Fig. 5 – Relative amount of defects computed by MD in Al at 10 K (Eth = 16 eV). FP: Frenkel pairs; Vf: free vacancies; Vc: vacancy clusters; If: free interstitials and Ic: interstitial clusters

DEFECT ENERGETICS

Formation and binding energies

The MDCASK code has been used to calculate the formation energy of the defect clusters. In these simulations a cluster is created with a defined size in a well relaxed matrix box, that is then further relaxed until the minimum energy state is reached. No calculation scheme among the numerous available in the literature is more accurate than the other (salesman problem). The one described by Soneda and de la Rubia [19] has been used in this work. The atoms that will form the desired cluster are created close to each other in the central region of a big enough box (84000 and 426000 atoms for vacancies and interstitials, respectively). The temperature of the system is increased artificially to about 0.8 Tm for an annealing period of 1 to 2 psec and then decreased step-wise so to take about 5 psec. to reach 200 K. To further reach 0 K, a quench scheme is used, which consists of setting the velocities of the atoms equal to zero each time the vector velocity and acceleration have opposite directions. This method has been used to determine the formation energies of clusters of up to 37 units in size for both interstitials and vacancies in Al. The results of such simulations are plotted in Figs. 5 for vacancies and 6 for interstitials. The investigation of the structure of the relaxed clusters will be reported elsewhere [34].

The binding energies of a defect in a cluster of size *(n)* are deduced from the formation energies E_f using the formula:

$$E_b^{\alpha}(n) = (E_f^{\alpha}(n-1) + E_f^{\alpha}(1)) - E_f^{\alpha}(n) \qquad (3)$$

where α is either *i* (interstitial cluster) or *v* (vacancy cluster). For larger clusters, the values of their binding energy are calculated from the extrapolation of the fits to those of smaller ones according to the expressions,
for vacancy clusters:

$$E_b^{v} = 0.61 - 0.8\,(n^{2/3} - (n-1)^{2/3}) \qquad (4)$$

for interstitial clusters:

$$E_b^{i} = 2.77 - 2.85\,(n^{2/3} - (n-1)^{2/3}) \qquad (5)$$

It is worth noticing at this point, that the interstitials in their clusters are more stable than vacancies.

Migration of point defects and small defect clusters

The diffusion simulation has be carried out using the same MDCASK code. The box matrix temperature is raised to the desired one (Langevin dynamics scheme) and the atoms are left to evolve without any external force, while recording the mean square displacements *(msd)* over all atoms in the computational cell containing the defect to be studied. It is important to notice that the number of jumps performed by the defect should be statistically large enough for the Einstein equation to be applicable:

$$D = \frac{msd}{6tn} \qquad (6)$$

where: $msd = \sum_i [r_i(t) - r_i(0)]^2$

and *t* is the simulation (diffusion) time and *n* is the defect size.

In the present work, the diffusion of single and di-interstitial were simulated in a (16x16x16) a_0^3 box at three different temperatures: 400, 600 and 800 K, for more than 500 psec. each. The diffusion of single and di-vacancy has

been simulated in the same manner. The diffusion single vacancy, however, was simulated only at 800 K, in order to check the compatibility of this simulation with the available experimental data. The calculated pre-factor and migration energies are listed in Table 1, together with the available experimental data [35]. It can been seen there that the results of the simulations are in good agreement with the experimental data, except for the di-vacancy migration energy, which is almost a factor two lower that the experimental one. It is known though, that experimentally it is difficult to obtain directly the migration energy of di-vacancy diffusion. It is usually deduced from the curvature of the Arrhenius plot of self-diffusion when this is measured over a wide temperature range [36].

Defect	$D_0/10^4 cm^2.s^{-1}$	E_m / eV	Experimental Data	
			$D_0/10^4 cm^2.s^{-1}$	E_m / eV
1V	440	062	470	0.61
2V	12.1	027	–	0.5
1I	15	012	–	0.11
2I	7.5	015	–	012
3I-9I	$D_0(2I)/n$	015	One dimensional diffusion mechanism	

Table 1 – Defect migration data in Al, calculated from MD

The diffusion of bigger clusters is not yet studied, but following the proposal of Trinkaus et al. [37] and Singh et al. [18], interstitial clusters containing between 3 and 9 interstitials are considered to be mobile. Their diffusion occurs via one-dimensional glide and the value of their migration energy is the same as that of the di-interstitials, but their pre-exponential factor is scaled to the size of the cluster as indicated in the last row of Table 1.

Defect interactions

Reaction	Product	Capture Radius (nm)
I+V	0	0.285
I+nV	(n-1)V	$1.15 \times 0.285 \times n^{1/3}$
I+nI	(n+1)I	$1.15 \times 0.285 \times n^{1/3}$
V+nV	(n+1)V	$0.285 \times n^{1/3}$
V+nI	(n-1)I	$0.285 \times n^{1/3}$

Table II – Defect interaction model

One of the most important ingredients in the study of damage evolution, is the definition of the capture radius of the interaction between different type of defects. If it is obvious for the interaction between single interstitial and vacancy, it is not well known neither experimentally nor theoretically for clusters of different type. Following the suggestions of Trinkaus et al [38], it is assumed that the capture radius of a single interstitial to be trapped by a cluster of either vacancy or interstitial type is 15 % bigger that the one for a single vacancy. The values used as input in the KMC simulations are listed in Table 2.

DAMAGE EVOLUTION

The results reported so far are based on the number of defects obtained at 10 K, after 10 to 20 psec. of the cascade event initiation. In order to calculate the fraction of defects that undergoes recombination in their nascent cascade or escape it at elevated temperature, the KMC method is used after implementing in it all the results discussed in the above sections.

Thermal annealing

In these simulations the starting microstructure produced by a given PKA energy is introduced in the centre of the KMC computational box of $(100)^3$ nm^3 in size. This microstructure consist of the spatial configuration of the defects obtained by the MD simulations of the corresponding displacement cascade. The system is then annealed at a given temperature for a given period of time, during which the type and the spatial and size distribution of the defects are recorded. These simulations have been performed for all the previously MD simulated displacement cascades, for temperature ranging from 10 to 700 K and for a period of time not less than 100 seconds. Figs. 6 to 9 show typical examples of the time evolution of the different defects. These example have been chosen on one hand to demonstrate the temperature effect and to assess any PKA energy effects present on the other. Four cases are analysed:

Low energy- low temperature

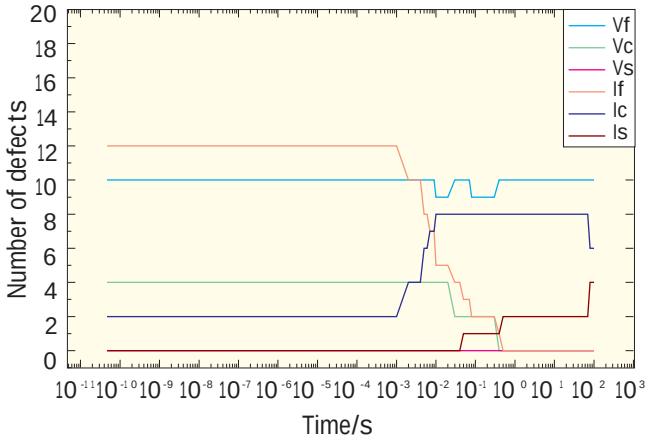

Fig. 6 – Isochronal annealing of the damage produced by a 3 keV PKA in Al at 75 K. Vf, If: free vacancies and interstitials; Vc, Ic: vacancy and interstitial clusters; . Vs, Is: vacancies and interstitials arriving at the box surface

This case is illustrated in fig 6, where the time evolution of the damage produced by a 3kev PKA energy displacement cascade (corresponding to fig. 2) is annealed at 75 K. It can

be seen that single interstitials (I_f), which were more than 70% of the interstitials present in the initial configuration, start to move and agglomerate after 0.001 sec. of anneal.

The clusters thus formed (I_c) are mobile and start reaching the simulation box surfaces at about 80 sec. Vacancies on the other hand, of which 12% are in clusters (V_c) at the initial stage (10 ps), start to annihilate due to the arrival of interstitials. The majority of the single vacancies (V_f) though, stay in this state even after 100 s. Additional simulations showed that their state does never change. It can be then concluded that at this temperature only interstitials are mobile, a strong indication that this temperature is in any case below stage III in the defect annealing terminology [39] and probably within stage I or II. This uncertainely will be clarified in the next sections.

Low energy-high temperature

The same spatial configuration as in the first case is taken as a starting distribution of the defects, i.e. the one produced by a 3 keV PKA, but it is annealed at room temperature (300 K). The temperature is defined as *high* since it is almost 1/3 of the meting temperature of Al. At this temperature, the atomistic mechanisms occurring during the damage evolution are very much different from case (a) as it is shown in fig. 7. In fact, all single interstitials disappear in less than 1 ms anneal, forming clusters in less than 1ns. The clustering mechanism of interstitials seems to be the dominating one between 1ns and 1 ms. Within 300 ms all interstitials escape the KMC simulation box which is 100 nm in edge. Vacancies that are in clusters at the initial stage, dissociate only after 100 psec. anneal and evaporate or annihilate completely after 0.01 sec, at which time the single vacancies start reaching the surface of the box massively. It can be then concluded that all types of defects are mobile at this temperature and that the vacancy clusters are not stable. These are the typical features of defects in stage III annealing.

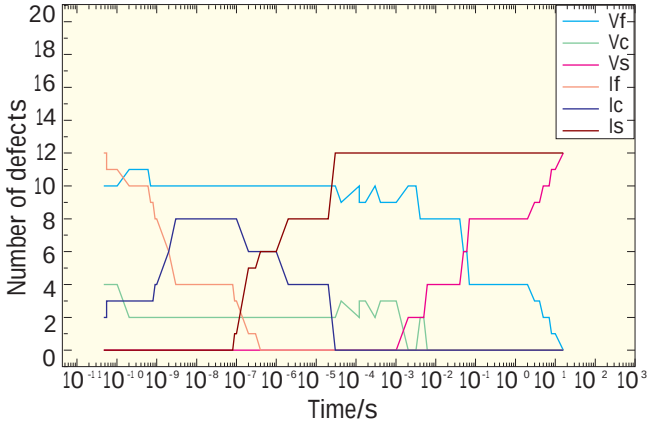

Fig. 7 – Isochronal annealing of the damage produced by a 3 keV PKA in Al at 300 K. Legend as in fig. 6

High energy-low temperature

The initial spatial configuration of the damage is taken to be the result of the MD simulation of a 10 keV PKA displacement cascade, shown in fig. 2. Basically the same mechanisms described in case (a) operate (see fig. 8), except that in this case there are more defects and they are more compactly distributed (fig. 2). The implication is that there is more annihilation of vacancies, either singles or clusters. But the vacancies are still immobile.

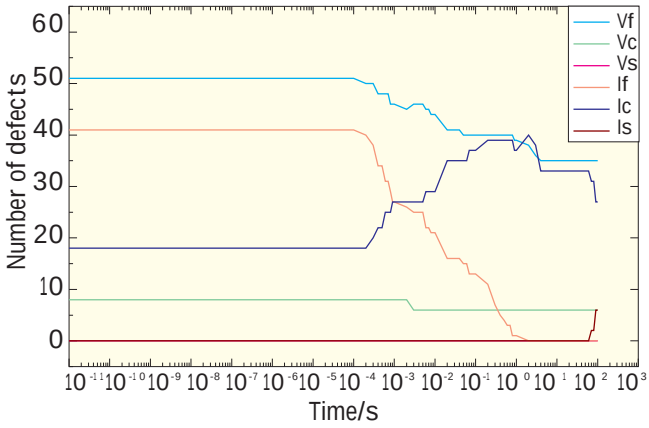

Fig.8 – Isochronal annealing of the damage produced by a 10 keV PKA in Al at 75 K. Legend as in fig. 6

High energy-high temperature

Here again, the same initial state as in the third case is taken. The mobility of vacancies promotes annihilation and escaping to the surface at a very early time for the same reason as for the second case. Previously though, all interstitials have escaped as can be seen in fig. 9.

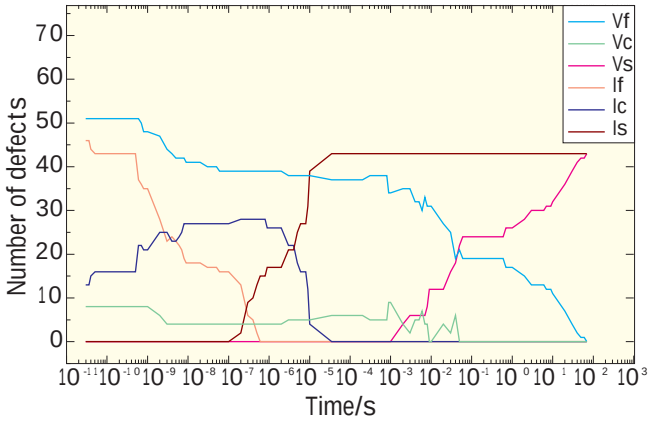

Fig. 9 – Isochronal annealing of the damage produced by a 10 keV PKA in Al at 300 K. Legend as in fig. 6

It is interesting to note here, that the vacancy clusters does not survive annealing even at temperatures as low as 75 K in contrast to other fcc metals e.g. Ni [28]. The other interesting fact that has to be noticed is the clustering of the interstitials prior to escaping. The analysis of the cluster size distribution during the evolution of the damage for all the PKA energies shows that the biggest interstitial cluster that might formed contains at most 8 interstitials. These observations may explain the fact that experimentally there are no quantified defect cluster observations. The results of this section present a further prove of the correctness of our calculated diffusion data and the binding energies of small clusters.

Freely Migrating defects

In the previous section, it has been shown that soon after their production, a non negligible amount of defects can escape their nascent cascade and migrate over long distances to reach the surface of the KMC box, if they are not trapped by other defects on their way. Those who escape are the so called freely migrating defects. It is obvious from the previous section that this amount will depend strongly on temperature and that moreover, this amount will contain different fractions of interstitials and vacancies. In this section, first the amount and the type of the defects that stay inside or escape the KMC box after thermal annealing of 100 sec. at temperatures ranging from 10 K to 750 K is determined. Secondly, the production efficiency of freely migrating defects, defined as the number of escaped defects normalised by the defect production value predicted by the NRT model, is evaluated at different temperatures. The productions efficiencies reported below are the average over 50 to 100 simulation runs of the same anneal for the same energy but different random number, in order to decrease the statistical discrepancies.

The number of the defects that stay single, cluster in, or escape the KMC box after a 100s annealing and normalised to the defect production value obtained by MD simulations, is plotted as a function of the annealing temperature in figs 10 and 11 for interstitials and vacancies, respectively. Only the low energy case will be analysed, as no fundamental difference has been found between the high and low energy cases in what concerns defect mobility, as discussed in the previous section.

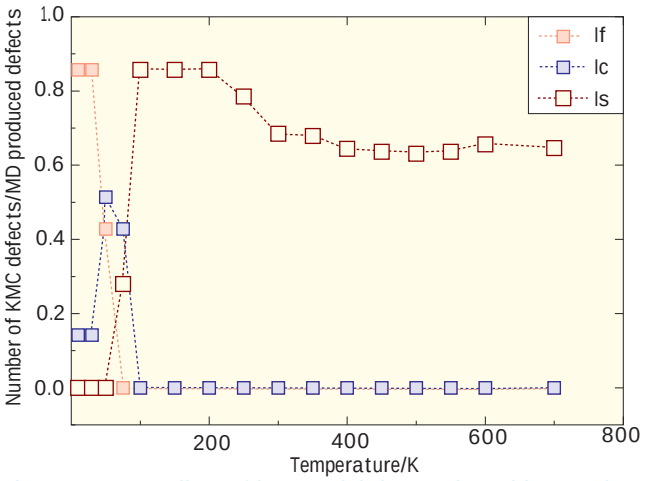

Fig. 10 – Annealing of interstitials produced by a 3 keV PKA in Al

Fig. 10 shows that the fraction of single interstitials that survive the thermal spike quench disappear as soon as the temperature is a little above 30 K. In other words, the single interstitials start moving at temperatures above about 30 K, which corresponds to the stage I of defect annealing in Al, in very good agreement with the experimental value of 37 K [35]. Just when the single interstitials start moving, a substage, where the interstitial clustering is promoted, appears. Experimentally, a substage has also been found at this temperature, but it has been interpreted as di-interstitials migration or stage II of defect annealing. At about 100 K, all types of interstitials are very mobile and, if they do not annihilate with the vacancies they may meet in their diffusion path, they reach the surface of the KMC box very fast.

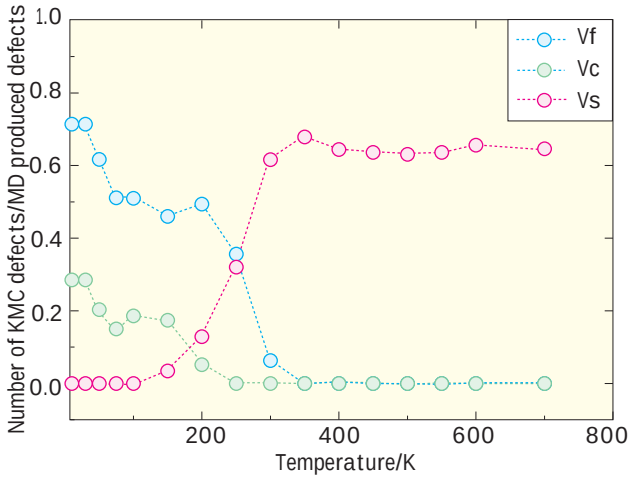

Fig. 11 – Annealing of vacancies produced by a 3 keV PKA in Al

Similarly, fig. 11 shows the same type fraction, but of vacancies. Here, at temperatures between 50 and 150 K, the fraction of single vacancies as well as that of cluster decreases slightly due to the annihilation with mobile interstitials. At about 200 K, all vacancy clusters are either annihilated or dissociated and the vacancies start to reach the surface, but their majority stays still in the box. Above 300 K all vacancies that survive annihilation reach the surface of the KMC box, i.e. the vacancies in their turn start to be highly mobile, which corresponds to the stage III of defect annealing.

At about 450 K a slight decrease of the fraction of the defects of either interstitial or vacancy type can be seen which might correspond to the stage V of defect annealing. However, experimentally this stage has been interpreted to correspond to the dissolution of large vacancy clusters. In the present computer simulation investigation only single cascades are considered and no large vacancy clusters have been found in any of the configurations issued from the displacement cascades simulated.

The production efficiencies are plotted as function of the energy of the MD simulated PKA in fig. 12, for three different temperatures. Each of the temperatures chosen is representative of the stages discussed above. At high temperatures, corresponding to temperatures above stage III, both types of defects are mobile and the same fraction of interstitials and vacancies escape the nascent cascade and reach the surface. This fraction decreases rapidly as the recoil energy increases and saturates at about 15 % of the NRT model predictions for PKA energies larger 3 keV. This result is comparable to the one obtained in Fe at 600 K, using the same simulation methods [4]. Approximately the same value has been estimated from irradiation enhanced diffusion experiments in Ni [40] and from diffuse X-ray scattering experiments following fission irradiation in Ni and Cu [41].

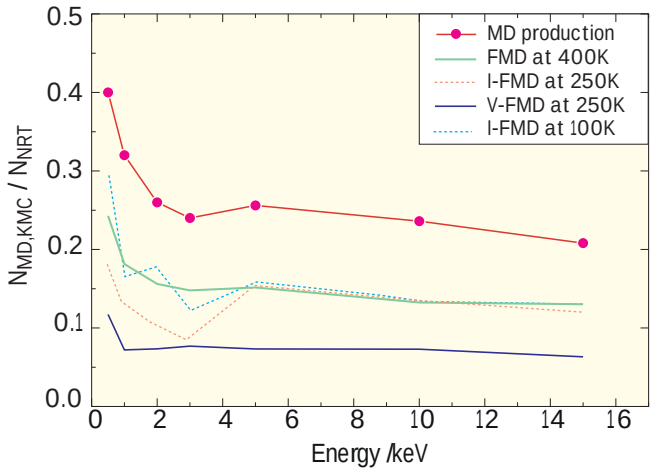

Fig. 12 – Freely migrating defects in Al at different temperatures

At 250 K, a temperature situated just above stage II but below stage III, where single interstitials are mobile, their clustering is more or less promoted but the vacancies are not yet very mobile and their clusters start to dissociate. From fig. 12, it can be seen that the fraction of interstitials that escape is much higher than that of vacancies. The interstitial fraction is about the same as the one obtained at 400 K for recoil energies higher that 3 keV. However, below 3 keV this fraction is very small, probably due to the fact that at these low energies the cascades are still not subdivided (more dense than at higher energies), which allows for much more intracascade defect annihilation. The fraction of freely migrating vacancies is rather low, at about 4%.

At much lower temperature, 100 K, only the interstitials are mobile and their freely migrating fraction is much the same as the one obtained at 400 K.

These results indicate that the determination of the freely migrating defect fraction, which is a key issue in predicting the microstructural evolution of the irradiation damage in metals, has to be related to the relative temperature where it has been calculated or measured.

Conclusions and Perspectives

A new hierachical computer model has been developed to simulate over a wide range of time and space the evolution of damage. From the present study the following conclusions can be derived:

- MD simulations of displacement cascades in Al show that there is little intracascade clustering and that those clusters formed are very small. The threshold for subcascade formation has been found to be around 5 keV, much smaller than for other fcc metals.
- MD methods have also been used to calculate the formation and migration energies of single defects and clusters. There is a good agreement with experimental values, where these are available.
- Kinetic Monte Carlo methods have been used in annealing simulations based on autoconsistent defect data and configurations produced by MD.
- The compatibility of these results with the experimental ones is a further test of the correctness of the calculated diffusion coefficients and binding energies.
- No significant difference is found between Al and other fcc metals in terms of defect behaviour (FMD fraction) above stage III.

In the next steps of the research program, it is expected to extend the model to:

- Account for the presence of the nuclear generated impurities, typically He and H, in order to deal more realistically with their possible stabilising effect on the defect population.
- Extend the model to more technologically interesting metals and alloys, in order to elucidate the atomistic mechanisms occurring during the lifetime of a technical material when subjected to intensive irradiation.

It is therefore our hope to make an underlying contribution to the definition of the materials needed to construct a fusion reactor as an ecologically safe and stable new energy source for the future generations.

Acknowledgements

This work was partially financed by the Fonds National Suisse de la Recherche Scientique. The assistance of the SIC in implementing the MD code in the Cray T3D is gratefully acknowledged. One of the authors (A.A) wish to acknowledge the valuable help provided by Dr. R. Schaublin (EPFL-CRPP-FMT), Dr. M. Alurralde (CNEA, Argentina), Dr. M. Spaczer (Paul Scherrer Institute) and E. Alonso (LLNL-USA) during the accomplishment of this work.

References

[1] A. Caro, M.Victoria in *Atomic scale calculations in materials science*, eds. J. Tersoff, D. Vanderbilt and V. Vitek; Mat. Res. Symp. **141** (1989) 19
[2] A. Caro and M. Victoria,;Phys. Rev. A **40** (1989) 2287
[3] A. Caro, S. Debiaggi and M. Victoria; Phys. Rev. B **41** (1990) 913
[4] S. Prönnecke, A. Caro, M. Victoria, T. Diaz de la Rubia and M. W. Guinan; Mater. Res. **6** (1991) 483
[5] S. Debiaggi and A. Caro; Phys. Rev. B **46** (1992) 7322
[6] F. Pedraza, D. Farkas and A. Caro, Rad. Effects and Defects in Solids **118** (1991) 143
[7] M. Alurralde, A. Caro and M. Victoria; Mater. Sci. Forum **97-99** (1992) 111
[8] S. Debiaggi and A. Caro; J. Phys. Cond. Mater. **4** (1992) 3905
[9] H. van Swygenhoven and A. Caro; Phys. Rev. Lett. **70** (1993) 2098
[10] M. Alurralde, A. Caro and M. Victoria; J. Mater. Res. **8** (1993) 449
[11] A. Caro, M. Victoria and R. S. Averback; J. Mater. Res. **5** (1990) 1409

[12] T. Diaz de la Rubia, A. Caro and M. Spaczér, Phys. Rev. B **47** (1993) 11483
[13] M. Spaczér, A. Caro, M. Victoria and T. Diaz de la Rubia; Phys. Rev. B **50** (1994-II) 13204
[14] M. Spaczér, A. Caro, M. Victoria and T. Diaz de la Rubia, J. Nucl. Mater. **212-215** (1994) 164
[15] M. Spaczer, A. Almazouzi, R. Schaublin and M. Victoria; Effects and Defects in Solids
[16] D. J. Bacon, , J. Nucl. Mater. **251**, (1997) 71
[17] H.L. Heinish and B.N. Singh, J. Nucl. Mater. **251**, (1997) 71.
[18] B.N. Singh, S.I. Golubov, H. Trinkaus, A. Serra, Yu.N. Osetsky, and A.V. Barashev, J. Nucl. Mater. **251**, (1997) 107.
[19] N. Soneda, and T. Diaz de la Rubia, Phil. Mag. In press.
[20] E. Alonso, M. Tang, and T. Diaz de la Rubia, Mater. Res. Soc. Symp. Proc. In press.
[21] M.J. Caturla, T. Diaz de la Rubia, and G.H. Gilmer, in: R.J. Culberstone et al. (Eds.), *Materials Synthesis and Processing Using ion beams*, MRS Symposia Proceedings No. 316, Materials Research Society, Pittsburgh, PA, 1994, p. 111.
[22] M.J. Caturla, T. Diaz de la Rubia, L. A. Marques and G. H. Gilmer, Phys. Rev. B, 54 (1996)16683
[23] L. M. Davies, U. Von. Estorff and S. Crutzen in the International Workshop of Decommisionned Material Collection and Testing of Structural Integrity Purposes, Mol-Belgium, 1995.
[24] T. Diaz de la Rubia, and M.W. Guinan, J. Nucl. Mater. **174**, (1990) 151.
[25] F. Ercolessi, and J.B. Adams, Europhys. Lett., **26**, (1994) 883.
[26] H.L. Heinisch, J. Nucl. Mater. **117**, (1983) 46.
[27] M. Jonsson, M.J. Caturla, and T. Diaz de la Rubia, Phys. Rev. B in press.
[28] A. Almazouzi, M.J. Caturla, T. Diaz de la Rubia, and M. Victoria, to be published.
[29] D.J. Bacon, A.F. Calder, F. Gao, V.G. Kapinos, and S.J. Wooding, Nucl. Instr. and Meth. in Phys. Res. B, **102**, (1995) 37.
[30] R.E. Stoller, G.R. Odette, and B.D. Wirth, J. Nucl. Mater. **251**, (1997) 49.
[31] A. Almazouzi, M. Alurralde and M. Victoria, to be published.
[32] M.J. Norgett, M.T. Robinson, and I.M. Torrens, Nucl. Eng. Design, **33**, (1975) 50.
[33] D.J. Bacon, and T. Diaz de la Rubia, J. Nucl. Mater. **216**, (1994) 275.
[34] A. Almzouzi, R. Schaublin, and M. Victoria, to be published.
[35] P. Ehrhart, Landolt Bornstein, New Series III/25, 219. A.
[36] A. Almazouzi, M.-P. Macht, V. Naundorf, and G. Neumann, phys. stat. sol.(a), **167**, (1998) 15.
[37] H. Trinkaus, B.N. Singh, and A.J.E. Foreman, J. Nucl. Mater. **199**, (1992) 5.
[38] H. Trinkaus, B.N. Singh and A.J.E. Foreman, J. Nucl. Mater. **251**, (1997) 172.
[39] H. Wollenberg, Physical Metallurgy, R.W. Cahn and P. Haasen Eds., North Holland, Fourth Edition, Vol. II, (1996) p. 1622.
[40] P. Fielitz, M.-P. Macht, V. Naundorf and H. Wollenberger, J. Nucl. Mater. **251**, (1997) 123.
[41] P. Erhardt and R.S. Averback, J. Phys. F, **14**, (1984) 1365.

SWISS-TX: FIRST EXPERIENCES ON THE T0 SYSTEM

RALF GRUBER, EPFL, SERVICE INFORMATIQUE CENTRAL, YVES DUBOIS-PÈLERIN, EPFL-DGM, LABORATOIRE DE MÉCANIQUE DES FLUIDES, AND SWISS-TX GROUP

Les deux premières machines du projet Swiss-Tx ont été installées au centre de calcul de l'EPFL. Le Swiss-T0 ayant comme système d'exploitation Digital Unix a été accepté en mai 98 et le Swiss-T0(Dual) ayant comme système d'exploitation Windows NT vient d'être installé le 2 octobre. Les prochains pas -Swiss-T1 et Swiss-T2 - sont presentés ainsi que le nouveau switch 12x12 qui reliera les boxes de Compaq/Digital. Les premiers résultats avec le programme Speculoos utilisant la méthode spectrale sont discutés. En particulier, le bus de communication EasyNet de Supercomuting Systems montre une meilleure scalabilité que le bus Fast Ethernet.

The first two machines of the Swiss-Tx project have been installed at the Computing Centre of EPFL. The Swiss-T0 machine running Digital Unix has been accepted in May 98 and the Swiss-T0(Dual) machine running Windows NT has been installed on October 2nd 98. The next steps - Swiss-T1 and Swiss-T2 - are presented as well as the new 12x12 crossbar switch that will connect the computer boxes from Compaq/Digital. First results made with the CFD spectral program Speculoos on the Swiss-T0 are discussed. Specifically, the EasyNet communication bus shows a better scalability behaviour than the Fast Ethernet bus.

INTRODUCTION

Since 10 years, computer scientists have proposed workstation and PC clusters as an alternative to big systems for High Performance Computing. Low price, high performance workstations or PCs are connected by an existing but slow, high latency Fast Ethernet switch. Special resource management systems such as Condor, or PBS have been developed and are offered as a freeware to the research community, others like LSF, or Codine, are licensed. More integrated workstation and PC clusters were built using faster communication networks like Myrinet or Gigabit Ethernet. Typical representatives of such new low price massively parallel machines are Beowulf, NOW and SuperCluster, all three being very powerful prototype machines mostly used for specific applications.

More integrated commodity supercomputers, that have a single-machine look, are presently developed at SRC, Sandia National Laboratories and EPFL. The SRC machine consists of a unique switch that interconnects Pentium Pros and later Merced processors. A first prototype will be installed at ORNL (Oak Ridge National Laboratories) beginning 99. The Sandia machine is based on 128 Alpha processors linked together with Myrinet switches. The EPFL machines are built in a cooperation with Supercomputing Systems in Zurich, Compaq/Digital, ETHZ and CSCS and are described below in detail. The SOS (Sandia/Oak Ridge/Swiss) research cooperation aims at shaping and testing the most promising supercomputer trends. In a previous paper (See EPFL Supercomputing Review n.10, 1997), we have presented the EPFL project, the Swiss-Tx. This communication gives the latest status of the systems development and the results of the first production runs for one scientific code, Speculoos.

THE SWISS-TX COMMODITY SUPERCOMPUTER PROJECT

THE PARTNERSHIPS

Highly evolved HPC relevant research projects have been conducted in Switzerland during the last 20 years. In hardware, the Institut für Elektronik (IFE, Prof. A. Gunzinger) at ETHZ and the Supercomputing Systems (SCS) company have developed the EasyNet concept, and have built two supercomputers, called Music and Gigabooster, both being commercialised through SCS. At EPFL, the PATP (Parallel Application Technology Program) project in cooperation with four major American research institutions and Cray Research had as goal the development of high end massively parallel scalable applications enabling to solve new advanced problems. The staff of the EPFL computing centre has accumulated a great experience in running supercomputers, the LITH-GRIP (Laboratoire d'Informatique Théorique, Groupe d'Intérêt du Parallélisme) has gained a lot of valuable experience in parallel programming and the LSP (Laboratoire de Systèmes Périphériques) has realised distributed file systems that are used in practice. The laboratory of Prof. J. Maddocks at DMA-EPFL has developed systems in computational steering and in visualizing parameter evolution to control execution on parallel machines. An important project between NEC and CSCS/SCSC led to a programming environment for massively parallel machines.

All this knowhow has been put together to realise with Compaq/Digital this Swiss-Tx project that aims at developping the necessary know-how in key areas, essential to the construction of the Swiss-Tx series and to its usage. Precisely, great efforts are made in testing operating systems, compilers and application libraries, in scheduling (partitioning), in user statistics and accounting, in development of the 12x12 crossbar switch for Swiss-T1 and Swiss-T2, in development of the corresponding PCI adapter, in implementation of the communication libraries (FCI, MPI-lite, MPI, and virtual shared memory programming model), in testing and evaluating the prototype machines (benchmarking), in porting/optimization of test programs in science, business and economy, in programming tools (monitors, debuggers, analysers), in a parallel file system and I/O, in distributed archiving, and in computational steering and visualisation. It is planned to commercialise the Swiss-Tx concept, in particular the crossbar switch that will be described later in more detail.

THE MACHINES

Machine	T0	T0(Dual)	T1	T2
Date	Dec. 98	Sept. 98	1stQ99	1stQ00
#P	8	16	72	504
Peak Gflop/s	8	16	72	1008
Memory GBytes	2	8	36	252
Disk GBytes	64	170	800	5000
Archive TBytes	1	-	1	7
Operating system	DEC Unix	W NT	DEC Unix	not decided
Communi-cation system	EasyNet bus	EasyNet bus	12x12 crossbar	12x12 crossbar

Within the two year Swiss-Tx programme, it is foreseen to build altogether four parallel computers [1]. The first two will have eight boxes with one and two processors respectively. For 1stQ 99, it is envisaged to build a 32 dual processor boxes Swiss-T1 machine delivering 72 Gflop/s peak performance. With the Swiss-T2 machine, it will be

possible to reach one Tflop/s one year later. All these machines will be tested and evaluated by user applications that already efficiently run on parallel supercomputers.

The first machine, the Swiss-T0 has eight DIGITAL-Alpha 21164 processors with 8 Gflop/s peak performance, 2 GByte of main memory, 64 GByte of disk space, and a one TByte archiving system. The processors are connected through an EasyNet bus from SCS. This system has been installed at the computing centre of EPFL at the end of 1997. It runs under DIGITAL UNIX. F77, F90, C, C++, HPF, and the ScaLAPACK and BLAS 1-3 libraries are available over the Fast Ethernet. The message passing library MPI-lite (a subset of MPI) has been installed and executed through the hardware interpreted Fast Communication Interface (FCI) that enables a direct store from one processor into the memory of another processor. The entire MPI communication library will be developped on top of MPI-lite and FCI and will be ready for the Swiss-T1 machine.

First measurements show a bandwidth of 35 MByte/s for EasyNet and 10 MByte/s for Fast Ethernet and a latency of 5µs for FCI and 12µs for MPI-lite when communicating over EasyNet and 500 µs for MPICH when communicating over the Fast Ethernet.

The first Swiss prototype supercomputer Swiss-T1 (see figs 1 and 2) with 72 Gflop/s peak performance will be available in 1stQ 1999. This machine will consist of 6 processing nodes of 6 dual processor boxes, altogether 72 production processors, and of one frontend node that takes care of the resource management and of all the external interactions. It will have 36 GByte of main memory, 800 GByte of disk space, and 1TByte of archiving space. A 12x12 crossbar connects the six boxes in a node, and fully interconnects the seven nodes. This machine will include the new Digital Alpha 21264 processors with an eight times higher memory bandwidth and a performance increase of over a factor of two.

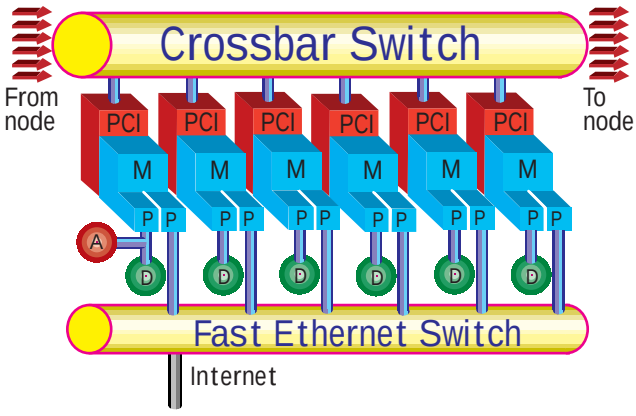

Fig. 1 – Swiss-T1 architecture is based on Alpha EV-6 processors. A node will consist of 6 dual processor boxes. They will be connected by a full 12x12 crossb`° switch based on the EasyNet concept. The remaining links are used to interconnect the 6 nodes. There will also be a Fast Ethernet. The users enter through the frontend which is fully integrated in the K-ring as a seventh node.

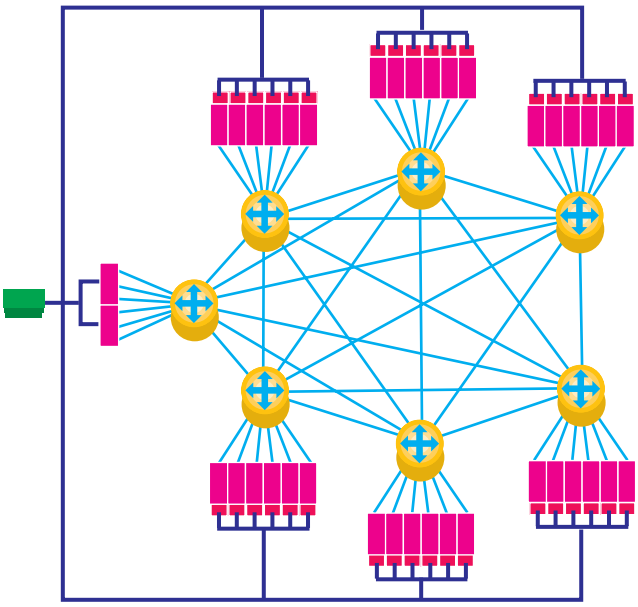

Fig.2 – The full connection between the 6 computational nodes and the integrated frontend at the left. The users will enter through the frontend computers

The Swiss-T2 machine will be based on the same architecture as the Swiss-T1. It is expected that the peak performance per processor and the number of processors per box will double, altogether a factor of 4 increase is predicted. With 504 processors, the peak performance will be over one Tflop/s (see fig. 3). It is anticipated that the machine could be ready by early 2000.

THE 12X12 CROSSBAR SWITCH AND THE K-RING ARCHITECTURE

The major advantages of the 12x12 crossbar switch are:

- fully non-blocking
- highly reliable
- communication bandwidth is presently 12x100 Mbyte/s, evolving to 12x400 Mbyte/s by 2000
- latency expected to be 5 µs for FCI and <20 µs for MPI
- excellent equilibrium by locally linking six boxes and six external links
- totally modular architecture: 12x12, 10x10, 8x8, 6x6 and 4x4 crossbars are possible
- operating system independent
- computer manufacturer independent
- all architectures are possible: we encourage the most optimal K-ring configuration

The K-ring architecture is shown in fig. 3. The 21 nodes are interconnected by three rings (K=3) in which jumps are made across 1, 5 and 8 crossbars. This gives the possibility to reach any node within at most two jumps. The routing table is constructed by a special program defining the optimal K-ring. There is also imbedded a control unit that

checks if the connections are realised according to the routing tables.

If N boxes have to be interconnected, a NxN switch can be developed. In the case of fig. 3, N=126, and the presented configuration is in fact a 126x126 switch. Such a switch has two level of communication speeds. In the first level the six local boxes are interconnected with a bandwidth of 100Mbyte/s per box. The second level of interconnection of the switches offers per box a smaller bandwidth. Typically, in the fig. 3 case each box has a 20 to 40 times smaller bandwidth. However, if we compare this bandwidth with the 35 Mbyte/s for the eight boxes measured on the Swiss-T0 (4 Mbyte/s per box), the long range bandwidth is comparable with the local EasyNet bandwidth. We shall see later that EasyNet is already very powerful for a CFD program.

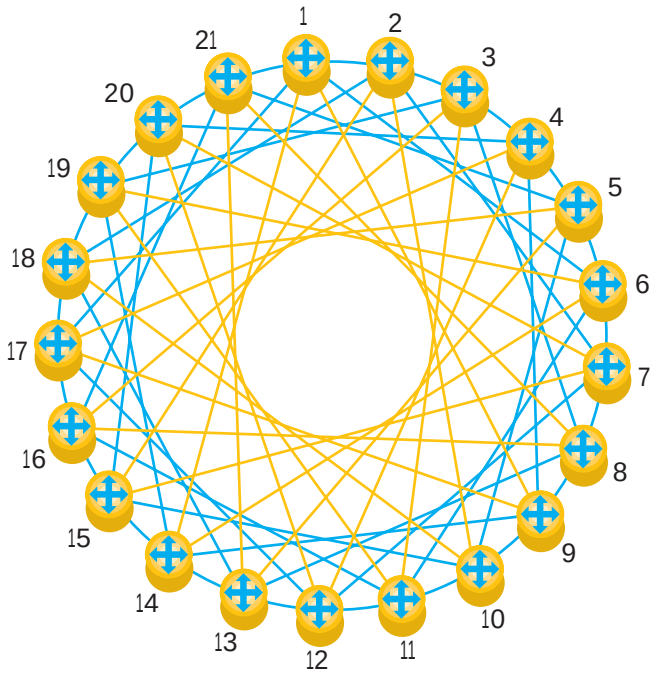

Fig. 3 – Swiss-T2 architecture with 21 (504 processors) nodes connected by a K-ring (K=3). Each of the vertices is a node with six quadri processor boxes connected by a crossbar.

SPECULOOS ON THE SWISS-TO

The parallel version of the Speculoos program [2] has been ported to the Swiss-TO to test performance of this prototype machine with four test cases. In Speculoos different types of PDEs are solved by the spectral element method based on the Ritz-Galerkin approach. In a mesh cell, the solution is approximated by a set of basis functions represented by l x m x n order polynomials in the three directions.

The program is written in object oriented manner using C++. The computation intensive parts are organised in such a way that highly efficient BLAS routines can be used. The options –O5 and –oldcxx are set for the compilation. Two versions of the program have been prepared, one using the MPICH communication library running over the Fast Ethernet and the FCI communication library running over the EasyNetBus. The two versions only vary by the calls to these two libraries. FCI cannot be used on one single processor. All the measurements are based on wall clock times.

TEST 1: 3 D POISSON EQUATION

No. processors	Elements per processor	MPICH speedup	FCI speedup
1	12 x 8 x 8	1	–
2	6 x 8 x 8	1.7	2.1
4	6 x 4 x 8	3.4	4
8	6 x 4 x 4	6.6	6.9

Table 1 – 3 D Poisson test case

The 3D Poisson equation is solved on a cubic domain discretized by 12 x 8 x 8 cells. In each cell basis functions with degrees 8, 10 and 11 in x, y and z, have been chosen. The resulting system of linear equations has about 700'000 unknowns. The system is solved by a conjugate gradient (CG) method, asking a tolerance of 10^{-14}, demanding 625 iteration steps. The single processor time is 2200 seconds. The speedup values for 2, 4 and 8 processors are shown in Table 1. In both cases, MPICH and FCI show a reasonable scaling.

TEST 2: STATIONARY 3 D STOKES PROBLEM

No. processors	Elements per processor	MPICH speedup	FCI speedup
1	4 x 4 x 4	1	–
2	2 x 4 x 4	1.7	1.9
4	2 x 2 x 4	2.4	3.3
8	2 x 2 x 2	3.0	5.8

Table 2 – 3 D Stokes problem

The domain consists of 4 x 4 x 4 intervals. The degrees of the polynomials are 5 in all three directions, leading to about 28'000 unknowns for the internal Helmholtz problem and 9300 unknowns for the external problem for the pressure. The pressure is computed by unpreconditioned CG method needing 92 iteration steps. The Helmholtz system is solved by a diagonally preconditioned CG method needing 125 iteration steps. Table 2 shows the speedup

measurements for MPICH over Fast Ethernet and FCI over EasyNet. The 8 processor speedup is almost two times higher for FCI than for MPICH. The single processor wall clock time was 2300 seconds.

TEST 3: 2 D HEAT EQUATION

No. processors	Elements per processor	MPICH speedup	FCI speedup
1	32 x 24	1	–
2	16 x 24	1.8	1.9
4	16 x 12	3.2	4
8	8 x 12	5.4	7

Table 3 – 2 D heat equation

The 2 D time evolutionary heat equation is solved on a non rectangular domain discretized in 32 x 24 quadrilateral cells. In each cell the basis function consists of a 4 degree polynomial in direction x and a 5 degree polynomial in the direction y, altogether 15'000 unknowns. The time integration is performed by a back difference formula (BDF) of third order. Ten time steps are made. Again a diagonally preconditioned CG method with a 10^{-14} tolerance needing 821 iterations has been chosen to solve the linear problems. Table 3 represents the scalability as a function of the number of processors for MPICH and FCI. The single processor wall clock time was 2300 seconds.

TEST 4: 2 D INCOMPRESSIBLE NAVIER-STOKES PROBLEM

No. processors	Elements per processor	MPICH speedup	FCI speedup
1	12 x 16	1	–
2	12 x 8	1.7	2.3
4	6 x 8	2.8	4.2
8	6 x 4	5.1	7.9

Table 4 – 2 D incompressible Navier-Stokes problem

The domain is a square cavity in which the fluid is moved by moving the cover. The velocity-pressure decoupling is performed by a LU decomposition generalised by Blair Perot and a pressure correction (BP1-PC). The time integration is realised using a BDF2 method and an extrapolative EX2 scheme for the non-linear term. Ten time steps are computed. The domain consists of 12 x 16 intervals, the polynomial degrees are 9 and 7 in the two directions, corresponding to about 12'000 points. The Reynolds number is 1000. Diagonal preconditioned CG are used to solve the Helmholtz system in 12 iteration steps. The pseudo-Poisson equation for the pressure has been solved by a P4 (Couzy) preconditioned CG method converging after 313 (resp. 471) iteration steps for the 1st (resp. 10th) time step. The single processor case takes 590 seconds.

CONCLUSION

First results achieved with the Speculoos program on the first Swiss-Tx prototype machine have shown the great advantage in performance of the EasyNet bus with respect to the Fast Ethernet.

With the 12 x 12 crossbar switch, the local interconnection between six boxes will have about 20 times higher bandwidth per box than the present EasyNet bus. The external communication bandwidth between the crossbars counted per box is of the order of what EasyNet can provide now on the T0 prototype machines. This means that real life applications like the object oriented spectral element program Speculoos are expected to perfectly scale up to the 504 processors Swiss-T2 machine with a 1 Tflop/s peak performance.

REFERENCES

[1] see ***http://capawww.epfl.ch/swiss-tx/***

[2] Y. Dubois-Pèlerin, V. Van Kemenade and M. Deville, *An Object-Oriented Toolbox for Spectral Element Analysis*, submitted to J. of Scientific Computing ■

EDITOR: Marie-Christine Sawley, SIC-EPFL
TEXT PROCESSING AND LAYOUT: Appoline Raposo de Barbosa, SIC-EPFL
ADDRESS: Service informatique central EPFL
MA-Ecublens Case Postale 121
CH - 1015 Lausanne
PHONE: (021) 693 22 11
FAX: (021) 693 22 20
E-MAIL: Marie-Christine.Sawley@epfl.ch
ADRESSE WEB: http://sawww.epfl.ch/SIC/SA/publications

RÉDACTEUR EN CHEF: Marie-Christine Sawley, SIC-EPFL
MISE EN PAGE ET GRAPHISME: Appoline Raposo de Barbosa, SIC-EPFL
ADRESSE: Service informatique central EPFL
MA-Ecublens Case Postale 121
CH - 1015 Lausanne
TÉLÉPHONE: (021) 693 22 11
TÉLÉCOPIE: (021) 693 22 20
ADRESSE ÉLECTRONIQUE: Marie-Christine.Sawley@epfl.ch
ADRESSE WEB: http://sawww.epfl.ch/SIC/SA/publications

America's Cup Yacht Design using Advanced Numerical Flow Simulations

Mario Caponnetto, Alessandro Castelli, EPFL–DGM – Hydraulic Machines Laboratory
Bernard Bonjour, Pierre-Louis Mathey, Stephane Sanchi, Mark L. Sawley,
EPFL–DGM – Fluid Mechanics Laboratory

La conception d'un voilier moderne de type 'Class America' s'appuie sur l'utilisation de simulations numériques de l'écoulement afin d'obtenir des prestations compétitives. Le calcul des écoulements hydrodynamique et aérodynamique autour de voiliers est, de toute manière, extremement exigeant; il requiert la technologie informatique et les techniques numériques les plus avangardistes. Un certain nombre de problèmes, critiques à la conception de voiliers pour l'America's Cup, sont discutés, et diverses approches liées à la simulation numérique avancée d'écoulements sont décrites.

The design of modern America's Cup racing yachts relies on the use of numerical flow simulations to obtain a competitive edge. The computation of the complex hydrodynamic and aerodynamic flows around sailing yachts is, however, extremely challenging, requiring state-of-the-art numerical techniques and computer technology. A number of the issues critical to America's Cup yacht design are discussed, and various approaches described to address them through advanced numerical flow simulation.

Introduction

The 30th America's Cup will be held in Auckland, New Zealand, commencing in October 1999 with the final races scheduled for February – March 2000. For the first time, there will be a Swiss competitor, the FAST2000 Challenge of the Club Nautique Morgien [1]. Participation in this prestigious sailing competition will pit local knowledge and skills against the world's sailing elite, and thus provide the potential to display local sailing experience and know-how to a world-wide forum.

FAST2000 has embarked on a two-yacht campaign, with the first yacht being launched at Pully on 27 August 1998. This yacht, which has been designed using the knowledge and techniques that existed during the last America's Cup in 1995, will be used for crew training. Three laboratories of the EPFL (LMF, LMH and LTC) are collaborating with FAST2000 in the design of the second yacht that will race in the Cup challenges [1].

The use of advanced techniques has become essential in the design of an International America's Cup Class (IACC) yacht for the America's Cup. Since a number of years, computational methods – in particular, numerical flow simulations – have been successfully applied to the design of sailing boats. Even though experimentation remains the tool most commonly used by designers to obtain accurate values of the hydrodynamic and aerodynamics forces acting on the boat, numerical simulations have some major advantages. In particular, they are relatively inexpensive and fast to use, so that it is possible to test and select different candidate geometries before setting up models for the towing tank or wind tunnel. Moreover, they allow the visualisation of several quantities – such as the flow streamlines, the wave profiles or the pressure distribution – that are very difficult to obtain from experiments. This is a very useful aid for the designer to understand the physics of the flow phenomena, at least from a qualitative point of view.

It is important to note that for IACC yachts, typical differences of speed between the winning and losing boats are about 1 – 2%. Under such conditions, it is clear that a high level of precision is required to predict boat performance to sufficient accuracy. Such precision places strong demands on both experimental and numerical methods used to determine the forces acting on the boat.

Most of the numerical simulations undertaken to date in this field have been based on potential flow theory, which reduces the complexity of the Navier-Stokes equations governing the flow and, consequently, the computational resources required. In particular, a large effort has been devoted to develop reliable tools (such as the panel method) for the computation of the wave resistance, as well as the lift and drag of appendages (keel and rudder) and sails. In some cases the basic hypothesis of the theory (inviscid flow) is satisfactory fulfilled, however, in a number of situations it has been shown that viscosity plays a fundamental role that can not be neglected. Nowadays, computational resources exist that allow the numerical simulation of the complex viscous flow around three-dimensional bodies, thus improving the accuracy of the computed flow solutions. While, to date, the computation-intensive nature of such simulations has severely limited their application to sailing boat design, these problems have been alleviated in the present study through the use of a high-performance computer system (Silicon Graphics Origin2000) installed at the EPFL.

The present paper briefly outlines some of the numerical simulation studies that are being undertaken at the EPFL within the framework of the collaboration with FAST2000.

Sailing Yacht Dynamics and Velocity Prediction

The forces acting on a sailing boat are the aerodynamic force applied to the sails, the hydrodynamic force applied to the hull and appendages and the force due to gravity [2]. For ideal steady motion, the sum of these forces and of their associated moments are both equal to zero. The sails develop a thrust and a lateral force that are respectively equal to the hydrodynamic resistance and lift generated by the hull, keel and rudder. Moreover, the aerodynamic lateral force and the hydrodynamic lift generate a heeling moment that must be compensated by the righting moment of the hull.

As a rule of the thumb, a boat is fast sailing upwind when it is stable (i.e., large and heavy and thus able to carry large sails even in strong wind) and when the sails and hull have a large efficiency (lift/drag ratio). On the contrary, when sailing downwind the boat should be light and narrow with the sails acting like parachutes to develop the maximum thrust. For the America's Cup, the race course is composed of a number of upwind and downwind legs. The fastest yacht exhibits the best compromise between upwind and downwind performances. In addition, each boat must comply with a class rule that imposes a relation between the boat length, displacement and sail area, with the goal of equalising the average speed of all the competing yachts. This rule, while representing a good average, is too simple to account in detail for the performances of different yachts in different weather conditions. The goal of each competitor is to take advantage of the simplistic nature of the rule, designing a yacht faster than that predicted by the rule. To achieve this, detailed knowledge of the forces acting on the boat is necessary. The approach used to predict the performance of a boat is to calculate (or measure) separately all the forces exerted on the hull and the sails as a function of the boat and wind speed, heel and yaw angles, and combinations of these variables. All these data are used by a computer code called a Velocity Prediction Program (VPP) to compute the final boat speed, and heel and yaw angles, that satisfy the equilibrium of the forces for a given wind speed and angle.

Hydrodynamic Forces

It is usual to consider the total drag of a boat sailing in calm water to be composed of wave, viscous and induced resistance. This convention has been adopted by naval architects for the purpose of towing tank experiments, however, it is not always easy to separate each component nor to take into account their mutual interaction. Sailing upwind in normal wind conditions, each of these three components represents approximately one third of the total resistance. Downwind, the induced resistance is negligible, the viscous resistance increases with the square of the boat speed, while the wave resistance has a much stronger dependence on boat speed (exponent of about 6 to 9). At high boat speed (strong winds), the wave resistance is predominant, while viscous resistance dominants at low boat speed (light winds).

Wave Resistance

Wave resistance is generated when the volume of the boat moves at the interface between the water and the air. The amount of wave resistance depends on the boat length and displacement, the shape of the hull and the Froude number (ratio of inertial and gravitational forces). While being essentially a potential flow phenomenon, some viscous boundary layer effects are present, especially close to the stern. The accurate numerical determination of wave resistance has always been extremely difficult for naval architects; for this reason, towing tanks are usually relied upon to determine the wave resistance of ships and yachts. From a computational point of view, the most efficient tools to solve potential flow problems are panel methods.

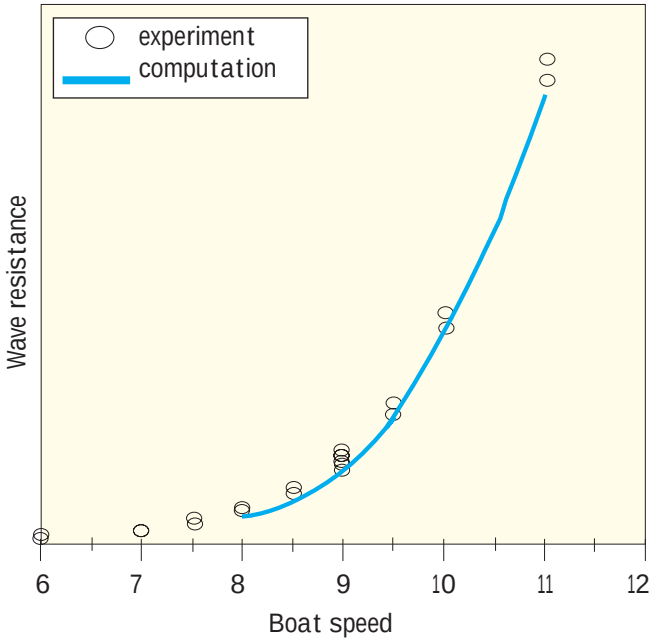

Fig.1 – Wave drag comparison between SHIPFLOW and experimental results

The commercial panel code SHIPFLOW developed by FLOWTECH International AB [3] is being intensively used at the EPFL to compute the free surface flow around yacht hulls at different speed and trim (yaw and heel angle). Rather that apply some approximative technique, SHIPFLOW uses an iterative procedure to satisfy the dynamics and boundary conditions on the exact water surface. During iteration, the trim and sink of the boat are adjusted to take properly into account the longitudinal moment generated by the towing force (applied at the centre of pressure of the sails) and the hull drag. These non-linear capabilities of SHIPFLOW are essential for the computation of the wave resistance of a sailing yacht. Indeed, one of the main problems for naval architects is to

predict how the *wetted length* of the boat changes when the exact profile of the wave and the trim are taken into account. The aim is obviously to design a boat shape that at speed has the greatest wetted length for the same length calculated statically by the above-mentioned class rule.

Previous attempts to calculate wave resistance of sailing boats with linear panel codes (where the boundary conditions are satisfied at the undisturbed flat-water surface) gave disappointing results. However, using the non-linear SHIPFLOW code, a good comparison between computations and the experimental data available has been obtained, as shown in fig. 1.

In addition to the quantitative determination of the wave resistance, visualisation of the computed flow fields allows an easy comparison of the pressure distribution and the wave profiles (fig. 2). Such information is very useful for the designer as an aid to understand the behaviour of different boat shapes from a qualitative point of view.

Fig. 2 – Waww profile around an IACC hull computed by SHIPFLOW

Viscous and Induced Resistance

Despite its importance, little attention is usually paid by naval architects to predict correctly the value of the viscous resistance. In the classical towing tank methodology, the friction drag is assumed to be equal to that of a flat plate having the same wetted area as the hull. Of course, the variable pressure distribution acting on the hull results in a behaviour of the boundary layer that is different than that of a flat plate. The ratio of the actual viscous resistance to that of the equivalent flat plate is called the "form factor", a quantity that is difficult to measure in a towing tank. It should be noted that a poorly-determined form factor can lead to an incorrectly-designed stern shape, which may give rise to undesired flow separation.

To determine the viscous resistance of an IACC yacht, three-dimensional Navier-Stokes computations for a number of different hull shapes are being performed at the EPFL using the commercial code FLUENT/UNS [4]. Each hull is usually tested at different speed and trim (fig. 3). The free surface is considered either as a flat undisturbed surface or using the perturbed free surface computed by SHIPFLOW. An accurate modelling of the boundary layer is required to calculate the skin friction and the occurrence of flow separation. A hybrid computational mesh is therefore adopted, using unstructured tetrahedral elements in the bulk flow, and structured prismatic elements in the boundary layer. Usually a total of 1.5 to 2 million cells is used, with about half of the cells located in the boundary layer. The computational time for the 500 iterations required for convergence is typically 10 hours on 12 processors of a Silicon Graphics Origin2000.

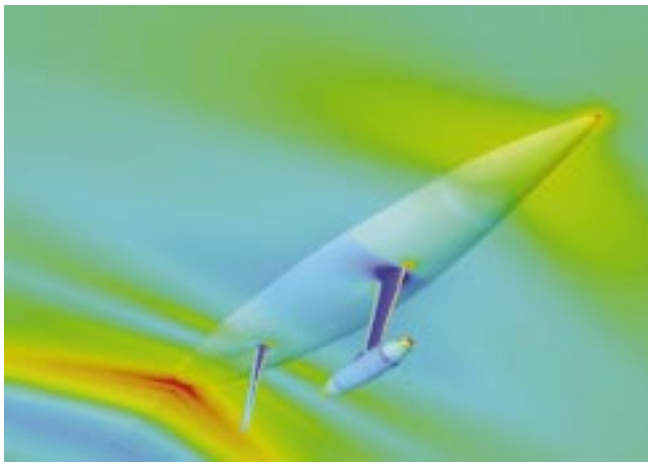

Fig. 3 – Pressure distribution around an IACC hull and appendages computed by FLUENT/UNS

When the boat sails with a yaw angle (angle of attack), the hull and the appendages develop a lift and an induced resistance that is determined by the vortex shedding. For wing-type surfaces, such as the keel and rudder, the location of the vortex separation can be easily predicted at the sharp trailing edge. A vortex smoothly detaches also on the leeward side of the hull, and in this case the location of the detachment and the strength of the vortex is strongly dependent on the pressure distribution along the hull and the behaviour of the boundary layer (see figure on cover page). Additional vortices are shed at the junction between the appendages and the hull (horseshoe vortex) that generate the so-called interference drag; this can be reduced with properly designed fillets and streaks. The strength and trajectory of all these vortices must be correctly calculated to predict the total lift and induced resistance of the boat and the optimum distribution of lift between the keel and the rudder.

Aerodynamic Flow

Sails are surfaces that must develop the maximum lift with the minimum drag when sailing upwind and the maximum drag downwind. Upwind they should work as

very efficient wings with the optimal circulation distribution and the lowest flow separation. Sails are not rigid wings, but thin flexible surfaces; this has some advantages and some drawbacks. The advantage is that the shape of the sail can be adjusted in the spanwise and chordwise directions by changing the tension of some wires (e.g. stays, runners) and by changing the flexion of the mast. The final shape of the sails is obtained as an equilibrium between the aerodynamic loads acting on the cloth and the internal and external loads in the cloth and the rig. The major drawback of having thin profiles is that the flow suddenly separates if the sail is not working at the ideal angle of attack.

Determination of Flying Shape of Sails

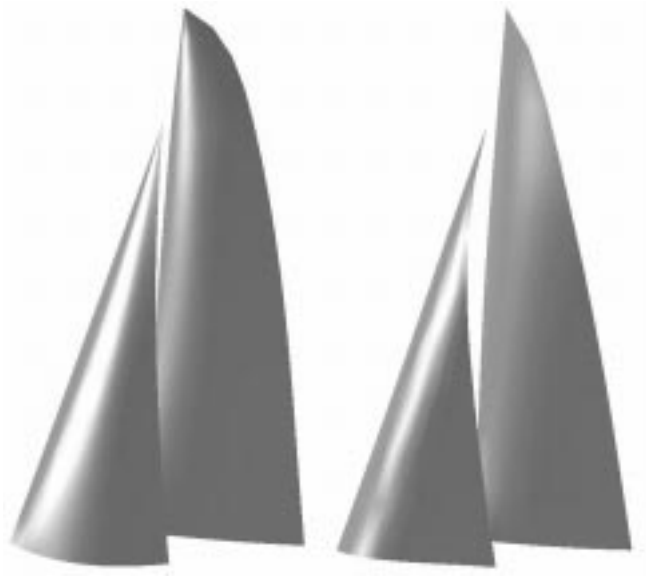

Fig.4 – Design (left) and flying (right) shapes of IACC upwind sails, the latter determined by FLOW-MEMBRAIN aeroelastic computations

For the design of optimum upwind sails, it is mandatory to be able to predict the deformed shape of the sails under aerodynamic loads (the so-called flying shape). The computation of such an aeroelastic phenomenon requires a coupling between flow and structural codes. This is performed at the EPFL using the codes FLOW and MEMBRAIN, developed by North Sails Inc. FLOW is a panel code for thin multiple lifting surfaces, that allows an accurate relaxation of the trailing wake. The input is the initial shape of the sails (without external loads) and the result is the aerodynamic pressure distribution corresponding to a given wind speed and direction. MEMBRAIN is a structural finite element code developed to calculate the stresses and the deformations of the sailcloth and rig under the action of a prescribed aerodynamic load distribution. FLOW and MEMBRAIN are used iteratively to compute the final sail shape and rig deformation satisfying simultaneously the aerodynamic and structural equilibrium. Fig. 4 shows the results of a computation, where the initial and flying shape of a sail plan can be compared. Even using a standard PC with a 300 MHz processor and 128 MB RAM, solving such an aeroelasticity problem with FLOW-MEMBRAIN is relatively simple and fast. It is therefore easy to compute a large number of different trims of the sails in a relatively short time (15-20 min. each), yielding the best-efficiency operating point of the sail.

Upwind Sail Computations

A potential flow computation using a panel code is able to give a precise aerodynamic load distribution only if the flow is attached. Under this hypotheses, the flying shape, lift and induced resistance are satisfactory predicted. However, even with well-trimmed sails some amount of separation always exists and this affects the final efficiency of the sails. In practice, it is important to know how the optimum camber of the sails varies with the wind speed, where the optimum is the best compromise between maximum lift achievable and minimum viscous resistance. Such effects can only be determined by a viscous flow computation. This can be performed with FLUENT/UNS, using the flying shape computed by FLOW-MEMBRAIN. Numerical simulations have been performed using a flow domain represented by about 1 million cells, with about 500 iterations needed to obtain convergence. A computational time of about 8 hours is required on 12 processors of the Silicon Graphics Origin2000.

One of the principal sources of aerodynamic drag is the mast, which can considerably increase the total aerodynamic resistance of the sail. Usually the flow separates on both sides of the mast and then reattaches on the mainsail (see fig.5). Lift is not affected in general, but the efficiency decreases. From the design point of view it is important to analyse this phenomenon in order to be able to design a mast profile that minimise the flow separation. A number of mast profiles have been studied using FLUENT/UNS. Fig. 6 shows a comparison between the computed pressure distribution and the experimental results obtained in the wind tunnel on a simple circular mast [5].

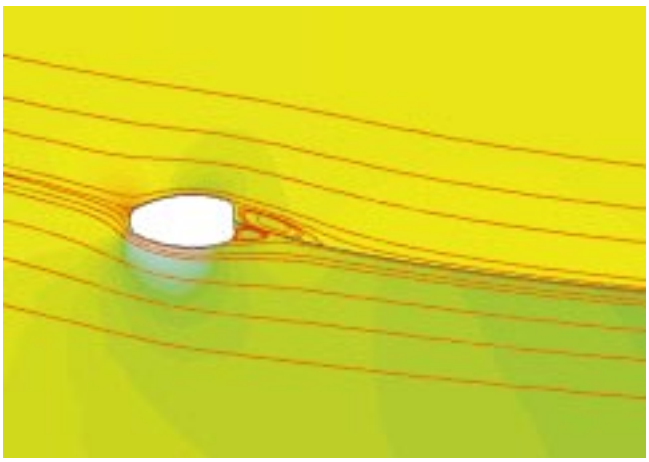

Fig. 5 – Flow around an IACC mast and mainsail geometry computed by FLUENT/UNS

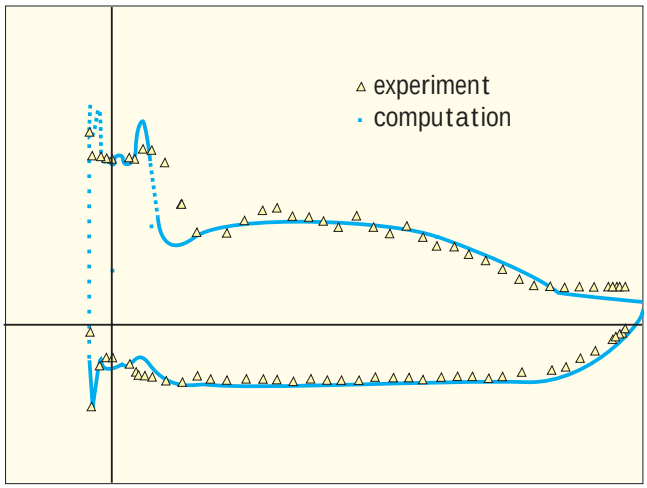

Fig. 6 – Comparison between FLUENT/UNS and experimental 2D pressure distribution around a mast and mainsail geometry

DOWNWIND SAIL COMPUTATIONS

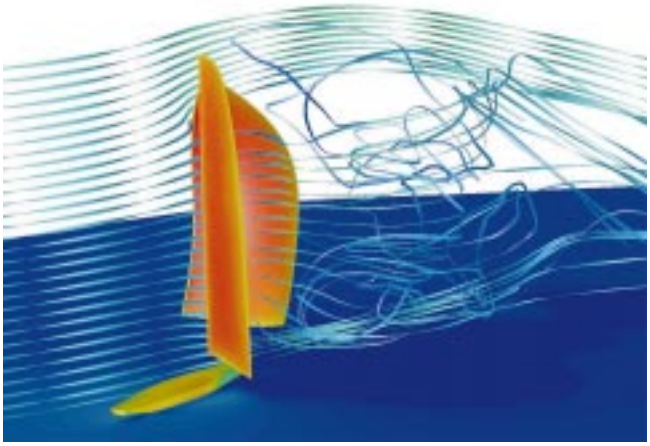

Fig. 7. – Surface pressure contours and streamlines for flow around a downwind sail configuration of an IACC yacht computed by FLUENT/UNS

While panel codes have been used with success to compute the forces acting on upwind sails, this is not possible for downwind sails due to the presence of large separated flow regions. A large potential for design improvement therefore exists in the use of advanced viscous flow codes for downwind flow configurations. Computations have been performed using FLUENT/UNS of the flow around such a downwind configuration, composed of a spinnaker and mainsail as well as the exposed section of the hull. (Such simulations required similar computational resources as for the viscous flow computations of the upwind configurations.) As shown in fig. 7, a large separated wake region is observed downwind of the spinnaker, comprised of a complex series of vortices. While it is not possible to compute accurately all the features of this extremely complicated flow behaviour, it is anticipated that such numerical simulations can provide a detailed knowledge of the interactions between the mainsail, spinnaker and hull. In addition, a study of the dependence of these interactions, as well as the general flow behaviour, for different sail configurations is of great interest for design purposes.

CONCLUSIONS

The present study has presented a number of different approaches that have been undertaken for the numerical simulation of the flow around a sailing yacht. It has been show that, despite the extremely complex flow behaviour present, information can be gained that is invaluable for design purposes. This information can be either quantitative evaluation of the hydrodynamic and aerodynamics forces exerted on the yacht, or detailed qualitative insights into the flow behaviour.

It can be recognised that the design of America's Cup yachts is entering a new phase through the use of state-of-the-art numerical flow solvers on high-performance parallel computer systems. These advanced tools provide valuable replacements for previously-employed empirical data and educated guessing. As a supplement to existing panel methods, they provide the naval architect with an enhanced probability of achieving a competitive design.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the intense interest in this work by members of the FAST 2000 Design Team (Philippe Briand, Sebastien Schmidt and Peter van Oossanen), with whom numerous discussions have led to a greater understanding of the challenges involved in sailing yacht analysis and design. This work has received financial support from the Commission pour la Technologie et l'Innovation.

REFERENCES

[1] M.-C. Sawley, FAST2000, un projet toutes voiles dehors, EPFL Supercomputing Review, 9, 18-20 (1997).

[2] J.H. Milgram, Fluid mechanics for sailing vessel design, Annual Review of Fluid Mechanics, 30, 613-653 (1998).

[3] Shipflow User's Manual; see also www.flowtech.com

[4] FLUENT/UNS User's Manual (Release 4.4), Fluent Inc. (1996); see also www.fluent.com

[5] S. Wilkinson, Ph.D Thesis, "Partially separated flow around masts and sails", University of Southampton, Ship Science Dpt., 1984. ■