

Grid Infrastructures for the Electronics Domain: Requirements and Early Prototypes from an EPSRC Pilot Project

R. Sinnott¹, A. Asenov², A. Brown², C. Millar^{1,2}, G. Roy², S. Roy², G. Stewart¹

¹National e-Science Centre, University of Glasgow

²Department of Electronics and Electrical Engineering, University of Glasgow

r.sinnott@dcs.gla.ac.uk

Abstract

The fundamental challenges facing future electronics design is to address the decreasing – atomistic - scale of transistor devices and to understand and predict the impact and statistical variability these have on design of circuits and systems. The EPSRC pilot project “Meeting the Design Challenges of nanoCMOS Electronics” (nanoCMOS) which began in October 2006 has been funded to explore this space. This paper outlines the key requirements that need to be addressed for Grid technology to support the various research strands in this domain, and shows early prototypes demonstrating how these requirements are being addressed.

1. Introduction

The relentless scaling of transistors in complementary metal oxide semiconductor (CMOS) integrated circuits has fuelled the phenomenal growth and success of the semiconductor industry, captured over the past 40 years by Moore’s law. Near the end of the current edition of the International Roadmap for Semiconductors (ITRS), in 2016, transistors will have reached sub-10 nm dimensions. The race between the major semiconductor manufacturers to demonstrate that they are capable of manufacturing devices with such dimensions has resulted in the demonstration of one-off MOSFETs with channel lengths of 15 nm by AMD in December 2000, 10 nm by Intel in June 2001 and 6 nm by IBM in December 2002 [1].

It is widely recognised however that intrinsic parameter fluctuations introduced by the discreteness of charge and matter is the major factor limiting the scaling and integration of devices as they approach nano-scale dimensions [2-6]. There are, at present no integrated simulation tools nor methodologies that can capture the full complexity of this problem and be used successfully to predict both the characteristics and scale of the intrinsic fluctuations in nanoCMOS transistors, and their subsequent impact on the performance of circuits and systems comprised thereof.

For example, variation in the number and position of dopant atoms makes each nano-transistor microscopically different, introducing significant parameter variations from device to device [7-15]. The trapping of a single electron in the channel region of a device can drastically change the current [16].

Furthermore interface roughness of the order of 1-2 atomic layers and related local variations in the oxide/body thickness introduce variations in gate tunnelling, quantum confinement and surface/bulk mobility from device to device [17-19]. The granularity of the gate material and the photo-resist both introduce unavoidable line edge roughness (LER) in the gate pattern definition and variations in geometry between devices [20-23].

The adoption of materials such as high-k gate dielectrics and SiGe introduce new fluctuations associated with random variations in their corresponding structure, composition, defects and strain [24]. When combined together, the variations in dopant/defect statistics, oxide/body thickness pattern, and gate material/geometry introduce intrinsic parameter fluctuations whose magnitude increases as devices shrink. These fluctuations have a crucial impact on the functionality, yield and reliability of circuits and systems at a time when fluctuation margins are shrinking due to continuing reduction in supply voltage and increased transistor count-per-chip [2,3].

In the presence of these intrinsic parameter fluctuations, the emphasis shifts from predicting the characteristics of a single nano-scale transistor to predicting the statistical behaviour of ensembles of macroscopically identical but microscopically different devices. 3-dimensional solutions with fine-grained discretisation are mandatory for studying atomic scale effects in nano-scale devices. Indeed the need for this statistical analysis transforms the problem into a four-dimensional one, where the fourth dimension is the size of the statistical sample.

Simulation complexity at the nano-scale is high due to the dominant role of quantum mechanical effects and complicated scattering mechanisms. Information from data and computationally intensive *ab initio* simulations supported by atomic-scale measurements are required to describe the atomic and electronic structure of the transistors. These have to be linked, at the analysis stage, back to the electrical behaviour of individual devices from the statistical sample. This statistical device analysis has to be followed by compact model

extraction for the whole statistical sample set, and statistical circuit simulation of an ensemble of circuits with identical topologies but microscopically different devices.

The EPSRC Pilot Project *Meeting the Design Challenges of nanoCMOS Electronics* (www.nanocmos.ac.uk) has been funded to explore and develop Grid based solutions addressing the research challenges inherent in this space. Through combining world leading electronics design research centres in the UK at the Universities of Glasgow, Edinburgh, Manchester, Southampton and York, supported with advanced Grid technologies we aim to revolutionise the electronics design industry. The project began in October 2006 and this paper captures the early experiences and software prototypes that have been developed, as well as the lessons learned thus far.

We note that the early focus of the project is targeted towards device modelling and developing Grid enabled atomistic services and compact models at the University of Glasgow. In the second year of the project, these services will provide the platform to support the higher level circuit and systems simulation requirements of the project. As a result, the experiences described in this paper are based around supporting the device modelling community at the University of Glasgow.

The rest of the paper is structured as follows. Section 2 gives an overview of the processes, services and associated data sets that the electronic design teams have identified as being crucial to support. Section 3 focuses on initial ideas on the design and development of this infrastructure in supporting the identified processes including how we are exploiting Shibboleth and advanced authorisation technologies to support the required fine grained data access to the electronics resources. Section 4 outlines the current implementation status and presents a snapshot of the existing prototypes. Finally in section 5 we conclude and provide a summary of our plans for the future.

2. Scientific Requirements

Development of a Grid infrastructure for the nanoCMOS electronics domain (or any other research domain) must address key requirements of that domain. To ensure that the Grid developers within the nanoCMOS project are developing technologies and solutions that address the scientific requirements, co-location of the scientific and Grid teams has taken place.

Numerous detailed discussions have taken place already within the context of the project

to better understand the process by which the scientists work, and the way in which they would like to work.

Diagrammatically, the overall sets of services and data are depicted in Figure 1. We expand on these services and data sets to the extent that they impact on the design and development of the Grid infrastructure. We note that for each of these services numerous instances using different technology bases can exist including both commercial offerings and academic software.




Figure 1: Conceptualisation of the nanoCMOS Services, Data and Design Processes

It is also the case that each of services (which we outline below) and data sets has strict intellectual property rights associated with them. These data sets themselves can be associated with particular technology nodes, e.g. 45nm scale devices from IBM; targeted towards particular types of device, e.g. low-power (LP) or high-power (HP); aimed at BSIM¹ or PSP² models.

The services and data sets themselves support the complete spectrum of processes associated with electronics design.

2.1 Process simulation

The process simulation is concerned with the physical steps necessary to turn a piece of silicon into a working device. This might include for example dopant implantation information; oxide growth; etching, deposition of metals etc. It is typically the case that this information is supplied by a commercial manufacturer such as IBM or Synopsis. This

¹ a widely adopted physics-based MOSFET model for circuit simulation and CMOS technology development.

² a compact surface-potential based MOSFET model modelling all relevant physical effects in present and upcoming bulk CMOS technologies.

information is provided typically as a *.tif* file. In future it may well be possible to generate these files directly through simulation. Currently the models explored in nanoCMOS assume this file is provided as an input.

2.2 Device simulation

A device simulation itself typically involves the solution of sets of coupled equations describing the distribution and flow of electrons in a given device structure. These might be based upon Drift Diffusion approaches solving Poisson's equation, or based on other approaches such as Monte Carlo based *ab initio* simulations.

The inputs to a device simulation will be information from the process simulation, i.e. the provided (or generated) *.tif* file, the dopant profile extracted for this device and further input file including information such as the simulation mesh used as the basis for the device simulation, and the oxide or nitride coating used with this particular device. Multiple (ensembles!) of device simulations need to be run to characterise the behaviour of a particular device. Each of these will have slightly different dopant profiles reflecting the variability inherent due to atomic scaling.

The output of a given device simulation is a current/voltage (I/V) curve describing the characteristics of that device with that particular dopant profile. The device simulation process as a whole generates many such I/V curves characterising the behaviour of that device with different dopant concentrations and distributions as shown in Figure 2.




Figure 2: Variations in the current-voltage characteristics of 200 transistors with different dopant distributions.

The device simulations themselves can be based on commercial tools such as TaurusTM or based on academic developed solutions. In the commercial case, license management is a crucial aspect to address for the Grid infrastructure.

The device simulations themselves can be extremely computationally intensive and generate extremely large data sets. As a result, a key requirement is to be able to access and

use large scale computational resources such as the National Grid Service (www.ngs.ac.uk) and ScotGrid (www.scotgrid.ac.uk).

The simulation software itself is primarily based upon Fortran 90 code.

2.3 Compact model extraction

Having generated the set of I/V curves for a particular device, it is necessary to abstract this information to a higher level so that multi-device circuit/system simulations can be performed. Compact models are semi-empirical analytical descriptions of the response of a device.

This process of generating a compact model is achieved through identification of an extraction strategy (identifying the subset of device model parameters which most influence the curves) and exploitation of commercial tools such as AuroraTM. It is the case that this phase will typically require domain knowledge and expertise in identifying the particular parameters that most influence the generated I/V curves.

This phase will typically require access to a larger shared memory type CPU resource.

2.4 Circuit simulation

Once compact models have been generated they can be used by circuit simulators to predict the behaviour of circuits and systems built from multiple combinations of these compact models. Typical examples of the kinds of behaviour analysed at the circuit/system level with these compact models are to identify how the set of connected components respond to a stepped input voltage or to explore particular tolerances of the integrated circuit.

Feedback at this stage can require modifications to the generated compact models which in turn may require device simulations to be redone. Linkage of device simulations, compact models and circuit/system simulations are an essential component of the nanoCMOS project in understanding how atomistic variation of devices impact upon system level design and simulation.

Both commercial and non-commercial applications are used for circuit simulation. One of the most widely used circuit simulators today is the Simulation Program with Integrated Circuit Emphasis (SPICE) simulator which has both open source and licensed versions.

3. nanoCMOS Grid Design

The nanoCMOS proposal recognised four key aspects that the Grid infrastructure must address. These include: security; data

management; workflows and resource management. At the time of writing work has focused primarily upon security and development of core simulation services supporting the processes identified in Figure 1.

It is the case that each of the services identified in Figure 1 requires simulation tasks to be completed. In the first year of the project we have decided to demonstrate at least one service for each of the simulation tasks: device simulation, compact model extraction and circuit simulation.

For the development of the Grid infrastructure the project has aligned itself with the OMII-UK technologies (www.omii.ac.uk). There were several reasons for this. Perhaps most importantly was the close synergies of the electronics design process, e.g. in supporting workflows capturing the way in which the different electronics activities need to be co-ordinated; advanced services for job submission; for data management etc. We fully recognise that the electronics design community will require extensions to the existing OMII-UK middleware however, e.g. in enhancements to address the security/considerations and IP protection of electronics design processes. Thus for example the existing WS-Security model of OMII-UK software does not support the fine grained authorisation definition and decision support required by nanoCMOS electronics researchers. Rather the WS-Security model of OMII-UK currently is focused upon message level security.

The data management and workflow aspects of the OMII-UK project are still at an early phase and work has focused initially upon the development of a family of OMII-UK services which support the device modelling and compact model generation phases of electronics design. Fundamental to all of the different phases of nanoCMOS work and design processes is support for enhanced and fine grained security.

3.1 nanoCMOS Security

The National e-Science Centre at the University of Glasgow have a body of expertise in supporting fine grained access to resources be they services or data sets across a variety of domains. We recognised that existing authentication-only approaches as typified by most larger-scale Grids built upon X509 based public key infrastructures are insufficient for the electronics domain. Use of an X509 digital certificate to ensure the identity of an individual is only a first starting point in building a secure infrastructure. Instead this domain demands finer grained

expression of security policies and their clear enforcement by providers of services or data.

Numerous technological solutions have been put forward looking towards providing various enhanced Grid security models and solutions such as CAS [25], PERMIS [26] and VOMS [27]. Examples of how these compare to one another is described in [28]. Recent developments in Grid standardization [29] and associated implementations [30] have shown, however, how finer grained models of security can be achieved supporting authorization closely integrated with Grid solutions.

One common way that fine grained security can be achieved is through Role Based Access Control (RBAC). Such systems allow for definition of roles which are typically associated with given privileges on a system and as such, are less susceptible to change than individual user identities. The roles themselves are assigned to subjects (users) by issuing them with an X.509 attribute certificate (AC) [31]. These roles and ACs can in turn be used to form the security policies for a given site. Systems such as PERMIS allow for the expression of digitally signed (and hence tamper proof) security policies based upon triplets comprised of <Role, Target, Action>. A local authority – the Source of Authority (SoA) will specify policies based upon institutional roles, institutional resources (targets) and actions that can be performed on those resources. Once defined, these policies can be used to ensure that only users with appropriate roles (privileges) can access certain services or data resources and perform certain actions.

It has been shown [32,33] how such infrastructures can be defined and used as the basis for limiting access to Grid resources and data sets. Such systems predominantly work at the local authorization level, i.e. the policies apply to the local site only. With Grid based inter-institutional VOs this model of security is not the norm and collective understanding of inter-institutional security infrastructures and a variety of attribute authorities is needed.

Supporting multiple attribute authorities is something that the Internet2 community has focused on explicitly in the Shibboleth architecture and protocols. The UK academic community is currently in the process of deploying Shibboleth technologies (<http://shibboleth.internet2.edu/>) to support local (existing) methods of authentication for remote login to resources. Through this model, sites are expected to trust remote security infrastructures for example in establishing the identity of users (authentication) and their associated privileges (authorization). To

support this, the Shibboleth architecture and associated protocols identify several key components that should be supported including federations, Identity Providers (aka origins), Service Providers (aka targets) and optionally Where Are You From (WAYF) services. Through these components, end users will have single usernames and passwords from their home institutions which will provide for seamless access to a range of resources at collaborating institutions and service providers. Local security policies at service provider sites can then be used to restrict (authorize) what resources authenticated users are allowed access to.

The typically scenario that arises with Shibboleth is as follows: a user attempts to access a Shibboleth protected service or Service Provider (SP) more generally (step 1 in Figure 3), they are typically redirected to a WAYF server (step 2) that asks the user to pick their home Identity Provider (IdP) (step 3) from a list of known and trusted sites.

After the user has picked their home site, their browser is redirected to their site's authentication server, e.g. an LDAP repository, and the user is invited to log in. After successful authentication, the attributes for that user are collected, e.g. from a local attribute authority (step 4) and sent back to the SP in a digitally signed SAML assertion with information that they have successfully authenticated at their IdP.

The SP then uses this information to decide whether it accepts the authentication assertion and that the provided attributes are valid and sufficient to provide access to the local resources it has available. That is, it makes an authorisation decision (step 6). Based on this the services that this user is authorised to access and use are made available through client portlets in the portal (step 7).




Figure 3: Shibboleth Access to Portal

This security model offers several direct benefits over existing PKI based Grid security approaches. Firstly, and most importantly it offers much finer grained security possibilities since the sites decide what attributes need to be provided to make an authorisation decision.

Any identity provider or attribute authority not providing the correct (signed) attributes will not be given access. Secondly, the end user does not have to manage their own X.509 certificate and memorise a strong private key password. Rather they are only required to log in to their own local identity provider. Thus the model for accessing Grid resources and non-Grid resources will (should!) be common across UK academia. Thirdly once a user has signed in to one service provider, they are able to seamlessly move to others without re-authentication (assuming the attributes and authentication meets the local authorisation requirements for those resource providers).

Within the UK a single federation (www.sdss.ac.uk) has been established and a small set of attributes based around the eduPerson class has been defined to support basic interoperability between identity and service providers. The nanoCMOS project (and most other Grid projects/virtual organisations) will of course want to be able to define their own security attributes and how they are used to define and enforce security decisions particular to their own virtual organisation. The NeSC DyVOSE project [34] developed a delegation issuing service [35] which supported the dynamic delegation of authority whereby virtual organisation specific roles could be delegated to local or remote trusted individuals and then be assigned to individuals at those institutions (or further delegated to other individuals/institutions depending upon the policy). Hence the definition and distribution of VO-specific attributes is readily supported.

One of the open challenges that the nanoCMOS project is facing right now with regard to Grid security is how to support protect federated resources. Thus for example, in the current design the focal point of the nanoCMOS virtual design foundry is a portal through which the various device simulations, designs and associated data sets are made available as shown in Figure 3. This portal in effect is the service provider which is protected by Shibboleth. However, whilst Shibboleth allows seamless access to a service provider, i.e. the portal, the challenge is in passing these security attributes to potentially remote services to make their own authorisation decisions. That is, the nanoCMOS partner sites will themselves want to make their own local authorisation decisions rather than leaving this entirely up to the portal itself. Thus the delegation of the authorisation decision making process to the portal will not be sufficient since sites/partners want to make autonomous decisions based on their own local

policies. Currently the authorisation policy is based on a “what you can see is what you can do” model where WS-security is used to ensure that the portlets and service communications are secure, i.e. if the appropriate attribute is not returned the client portlet will not be deployed and hence the end user cannot access and use the associated service.

There are several possibilities that are being explored to overcome these kinds of issues. For example, one *belt and braces* approach is for the authorisation at the portal level to be augmented service level authorisation. For example, each *service* can be protected by Shibboleth and access to this is then based upon the service provider seamlessly (from the user perspective) accessing the attribute authority to determine that the user requesting access has the appropriate authority, i.e. the right attributes required. It could be argued that an authorisation decision is not needed if the attribute is delivered to the portal itself, i.e. authorised access to a portlet indicates that authorised access to the service is ensured. However, we recognise that autonomy of sites is crucial to support in this domain. In the recently funded JISC project VPman (www.nesc.ac.uk/hub/projects/vpman) we are exploring the integration of the VOMS attribute authority with the PERMIS authorisation infrastructure. In this scenario, a nanoCMOS wide virtual organisation attribute authority is defined which can be used for the definition of nanoCMOS roles used by local service providers to authorise access to local resources for example, as opposed to each partner site having their own attribute authority specific to the nanoCMOS project.

4. Implementation

At the time of writing, early prototypes of the nanoCMOS services have been implemented and made available within a project portal protected by Shibboleth. The portal itself contains numerous components including a project wiki and of course access to the various services that have been prototyped thus far.

The front end access to the portal is depicted in Figure 4. We note that this portal displays the various attributes that have been released by the identity provider and attribute authority at the University of Glasgow. We also note that through the recently started OMII-SP project at NeSC (www.nesc.ac.uk/hub/projects/omii-sp) we have developed portlets which allow scoping of these attributes. Thus in this case, the only attributes that will be recognised by the portal

are those prefixed with *NanoCMOS*. Furthermore, this scoping allows the portal to be restricted to only accept attributes from known and trusted sources, e.g. the nanoCMOS partner sites or more restrictively, only from specific individuals at those sites. Tools supporting more advanced attribute release and attribute acceptance policies targeted to the needs of Grid based virtual organisations are essential for the successful uptake of Shibboleth by the e-Research community.




Figure 4: nanoCMOS Portal with Scoped Attributes

The initial nanoCMOS services themselves that have been prototyped thus far in nanoCMOS have included the first version of a device modelling simulator (called Geronimo – shown in Figure 5) and a Grid enabled version of a commercial application Aurora used for creation of compact models shown in Figure 6.




Figure 5: Geronimo Device Modelling Portlet




Figure 6: Aurora Compact Model Portlet

Both of these services have been developed using OMII-UK software and have exploited components such as GridSAM. The Geronimo service itself has been developed so that it is able to support job submission to a variety of resources. Currently it is able to submit to resources such as a Sun Grid Engine based cluster in the electronics department, the Condor pool at NeSC. Work is currently on-going to support job submission to the National Grid Service and the ScotGrid cluster at Glasgow.

We also note that with regard to the Aurora service, we have shown how it is possible to use Shibboleth technologies and delivery of attributes to address one of the key problems facing the Grid community, namely: license management. Thus, if a user has an Aurora license (and this is delivered to the portal) they are able to see and invoke this service. In the current implementation the Aurora service is deployed on a server at the NeSC in Glasgow, but a license checked out from the FlexLM server in the electronics department.

At the time of writing work is on-going in development of the SPICE service and associated portlet.

5. Conclusions

The electronics design industry is facing numerous challenges caused by the decreasing scale of transistors and the increase in device design flexibility. To overcome these challenges demands that scientists co-ordinate their efforts. The collaborative nature of the Grid provides a compelling model closely aligned with the requirements of this community. One of the key requirements which must be at the heart of any Grid solution in the electronics domain is fine grained security.

The early prototypes developed have shown how Shibboleth combined with virtual organisation specific attributes can restrict access to services and data to the appropriate privileged users.

The work is of course still in its early phase, and much remains to be done including enhancing the existing services to support simulation of larger ensembles of devices, and importantly for extensions to Grid know how such as security oriented workflow definition and enactment. The requirements for these extensions first and foremost, come from the needs of the electronics research community. The project has attempted to closely integrate the Grid developer efforts and electronics design processes through co-location of researchers.

6. References

1. B. Doris et al. *Extreme scaling with ultrathin Si channel MOSFETs*, IEDM 2002 Tech. Dig. p. 267 (2002).
2. H. P. Tuinhout, *Impact of parametric mismatch and fluctuations on performance and yield of deep-submicron CMOS technologies*, Proc. ESSDERC, pp.95-101, Florence, Italy, 2002.
3. D. J. Frank and Y. Taur, *Design considerations for CMOS near the limits of scaling*, Solid-State Electronics, vol. 46, pp 315-320 (2002).
4. K. Takeuchi, R. Koh and T. Mogami, *A study of the threshold voltage variation for ultra-small bulk and SOI CMOS*, IEEE Trans. Electron Dev, vol. 48, p. 1995 (2001).
5. T. Mizuno, J. Okamura and A. Toriumi, *Experimental study of threshold voltage fluctuation due to statistical variation of channel dopant number in MOSFET's*, IEEE Trans. Electron Devices, Vol. 41, pp. 2216-2221, (1994).
6. A. R. Brown, A. Asenov and J. R. Watling, *Intrinsic Fluctuations in Sub 10 nm Double-Gate MOSFETs Introduced by Discreteness of Charge and Matter*, IEEE Transaction on Nanotechnology, Vol. 1 pp. 195-200 (2002).
7. P.A. Stolk, F.P. Widdershoven, D.B.M. Klaassen, *Device modeling of statistical dopant fluctuations in MOS transistors*, Proc. SISPAD'97, pp. 153 - 156, 1997.
8. Wong H.-S. and Taur Y. *Three dimensional 'atomistic' simulation of discrete random dopant distribution effects in sub-0.1 μ m MOSFETs*, Proc. IEDM Dig. Tech. Papers., pp. 705-708, 1993
9. D.J Frank, Y. Taur, M. Jeong and H.-S. P. Wong, *Monte Carlo modelling of threshold variation due to dopant fluctuations*, 1999 Symposium on VLSI Technology Digital Technology Papers, pp 169-170, 1999.
10. D. Vasileska, W. J. Gross and D. K. Ferry, "Modeling of deep-submicrometer MOSFETs: random impurity effects, threshold voltage shifts and gate capacitance attenuation", Extended Abstracts IWEC-6, Osaka 1998, IEEE Cat. No. 98EX116, pp. 259-262.
11. A. Asenov, S. Kaya and A. R. Brown, *Intrinsic Parameter Fluctuations in Decanometre MOSFETs Introduced by Gate Line Edge Roughness*, IEEE Trans. Electron Dev, 2005.
12. A. Asenov, G. Slavcheva, A.R. Brown, J.H. Davies and S. Saini, *Quantum enhancement of the random dopant*

- induced threshold voltage fluctuations in sub 100 nm MOSFETs: A 3-D density-gradient simulation study*, IEEE Trans. Electron Dev., vol. 48, pp. 722-729 (2001).
13. A. Asenov and S. Saini, *Polysilicon Gate Enhancement of the Random Dopant Induced Threshold Voltage Fluctuations in Sub 100 nm MOSFETs with Tunnelling Oxide*, IEEE Trans. Electron Dev., Vol. 47, No. 4, pp 805-812 (2000).
 15. A. Asenov, *Random dopant induced threshold voltage lowering and fluctuations in sub 50 nm MOSFETs: A 3D 'atomistic' simulation study*, Nanotechnology, Vol.10, pp.153-158, 1999.
 16. A. R. Brown, J. R. Watling and A. Asenov, *A 3-D Atomistic Study of Archetypal Double Gate MOSFET Structures*, J. Computational Electronics, Volume 1, pp. 165-169 (2002).
 17. A. Asenov, S. Kaya and A. R. Brown, *Implications of Imperfect Interfaces and Edges in Ultra-small MOSFET Characteristics*, Phys. Stat. Sol (b) Vol. 233, No. 1, pp. 101-112 (2002).
 18. A. Asenov, S. Kaya, J. H. Davies, *Intrinsic Threshold Voltage Fluctuations in Decanano MOSFETs due to Local Oxide Thickness Variations*, IEEE Trans. Electron Dev., vol. 49, pp. 112-119, 2002.
 19. A. Asenov, S. Kaya, J. H. Davies and S. Saini, *Oxide Thickness Variation Induced Threshold Voltage Fluctuations in Decanano MOSFETs: A 3D Density Gradient Simulation Study*, Superlattices and Microstructures, Vol. 28, No. 5/6, pp. 507-515, (2000).
 20. D. Vasileska, W. J. Gross and D. K. Ferry, *Modeling of deep-submicrometer MOSFETs: random impurity effects, threshold voltage shifts and gate capacitance attenuation*, Extended Abstracts IWEC-6, Osaka 1998, IEEE Cat. No. 98EX116, pp. 259-262.
 21. P. Oldiges, Q. Lin, K. Pertillo, M. Sanchez, M. Jeong, and M. Hargrove, *Modelling line edge roughness effects in sub 100 nm gate length devices*, Proc. SISPAD 2000, p. 31, 2000.
 22. T. Linton, M. Giles and P. Packan, *The impact of line edge roughness on 100 nm device performance*, Proc. Silicon Nanoelectronics Workshop, p. 82, 1998.
 23. T. D. Linton, S. Yu and R. Shaheed, *3D modeling of fluctuation effects in highly scaled VLSI devices*, VLSI Design, Vol.13, pp.103-109, 2002.
 24. M. Ono et al. *Effect of metal concentration nonuniformity in gate dielectric silicates on propagation delay time of CMOS invertors*, Proc. SSDM 2002 Nagoya, Japan, p. 710 (2002).
 25. L. Pearlman, et al., *A Community Authorisation Service for Group Collaboration*, in Proceedings of the IEEE 3rd International Workshop on Policies for Distributed Systems and Networks. 2002.
 26. D.W. Chadwick, A. Otenko, *The PERMIS X.509 Role Based Privilege Management Infrastructure*, Future Generation Computer Systems, 936 (2002) 1-13, December 2002. Elsevier Science BV.
 27. R. Alfieri, et al, *Managing Dynamic User Communities in a Grid of Autonomous Resources*, CHEP 2003, La Jolla, San Diego, March, 2003;
 28. A.J. Stell, *Grid Security: An Evaluation of Authorisation Infrastructures for Grid Computing*, MSc Dissertation, University of Glasgow, 2004.
 29. V. Welch, F. Siebenlist, D. Chadwick, S. Meder, L. Pearlman, *Use of SAML for OGSA Authorization*, June 2004, <https://forge.gridforum.org/projects/ogsa-authz>
 30. D.W. Chadwick, *An Authorisation Interface for the Grid*, Proceedings of UK e-Science All Hands Meeting, September 2003, Nottingham, England.
 31. ITU-T Recommendation X.509 (2001) | ISO/IEC 9594-8: 2001, Information technology – Open Systems Interconnection – Public-Key and Attribute Certificate Frameworks
 32. R.O. Sinnott, A.J. Stell, D.W. Chadwick, O. Otenko, *Experiences of Applying Advanced Grid Authorisation Infrastructures*, Proceedings of European Grid Conference (EGC), June 2005, Amsterdam, Holland.
 33. R.O. Sinnott, A.J. Stell, J. Watt, *Comparison of Advanced Authorisation Infrastructures for Grid Computing*, Proceedings of International Conference on High Performance Computing Systems and Applications, May 2005, Guelph, Canada.
 34. J. Watt, R.O. Sinnott, J. Koetsier, A.J. Stell, *DyVOSE Project: Experiences in Applying Privilege Management Infrastructures*, UK e-Science All Hands Meeting, Nottingham UK, September 2006.
 35. R.O. Sinnott, J. Watt, D.W. Chadwick, J. Koetsier, O. Otenko, T.A. Nguyen, *Supporting Decentralized, Security focused Dynamic Virtual Organizations across the Grid*, 2nd IEEE International Conference on e-Science and Grid Computing, Amsterdam, December 2006.