UDC 621.391

J.N. DAVIES[1], P. COMERFORD[1], V. GROUT[1], N. RVACHOVA[2], O. KORKH[2]

[1] *Creative and Applied Research for the Digital Society (CARDS), Glyndŵr University, Wrexham, UK*
[2] *Poltava National Technical University, Poltava, Ukraine*

OPTIMIZATION PRINCIPLES FOR ELIMINATING ACCESS CONTROL LISTS WITHIN A DOMAIN

The infrastructure of large networks is broken down into areas that have a common security policy called a domain. Security within a domain is commonly implemented at all nodes however this can have a negative effect on performance since it introduces a delay associated with packet filtering. When Access Control Lists (ACLs) are used within a router for this purpose then a significant overhead is introduced associated with this process. It is likely that identical checks are made at multiple points within a domain prior to a packet reaching its destination therefore by eliminating ACLs within a domain by modifying the ingress/egress points with equivalent functionality an improvement in the overall performance can be obtained.

Keywords*: Routing Domain, Performance, Delay through Routers, Access Control List, ACL optimization, Off-line verification of ACLs, Firewalls, Inter-Firewall Optimization, IP packet filtering.*

Introduction

Modern computer networks are expected to provide reliable high performance end to end connectivity at any point in the world. They must also provide the ability to filter packets so that access to services is limited to trusted traffic defined in the security policy for the network. This must be achieved with a minimal delay without compromising the security policy. It can be a challenge for a network manager to meet these two conflicting requirements.

Most networks contain one or multiple connections to external networks e.g. Internet which is considered a great security risk. To mitigate this, trusted networks are created which perform stringent security checks on packets travelling across the network boundary in either direction. Such networks operate under a common security policy managed by a single authority and are known as domains. If network traffic is filtered at all ingress and egress points in the network then it should only contain traffic which is defined as trusted under the security policy (Fig. 1) [1].

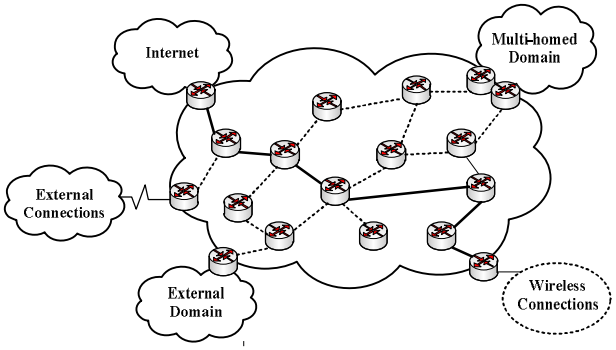

Fig. 1. Typical Domain allocation

Infrastructure security within a domain is normally implemented in either firewalls or routers containing Access Control Lists (ACLs). ACLs have a common implementation across all platforms e.g. Cisco, Juniper and Linux [2]. Since every packet has to be tested significant delays can result from the introduction of such techniques due to the filtering requirement [3]. Attempts have been made to use various techniques to optimize the delay through routers caused by ACLs [4]. Optimization of packet filtering performance has been the subject of intense research for the past decade [5]. A number of studies have identified rule relations within an ACL which may result in redundant or conflicting rules.

This paper investigates the significance of the delay encountered through the use of various ACL techniques. Factors which contribute to the delay incurred by packets passing through a router are identified and subsequently, a number of experiments were conducted to quantify these. Recommendations were made to give guidance to network engineers that can be used during the network design phase. An argument for the use of an Operating System (OS) with appropriate functionality for a given task is put forward based on experimental results. The authors recommend an optimal configuration using a worked example based on previous findings. Finally, a mechanism for the consolidation of distributed ACLs to a single ACL providing equivalent functionality is presented. Only delays through network equipment were considered in this paper.

1. Packet delays within a Router

When considering the packet delay through a domain there are a number of factors that need to be con-

© J.N. Davies, P. Comerford, V. Grout, N. Rvachova, O. Korkh

sidered. These factors include the route selected by the routing protocol, the bandwidth of the links along the selected route and the internal delays within the equipment [6]. Routing Protocols optimize the route selection using a shortest path algorithm based on cost functions for each path [7]. The delays experienced within equipment e.g. routers and switches are often ignored since the link bandwidth has generally been considered as the dominant factor [8]. However as technology has improved the link speeds have increased and so the equipment delays have become more significant [9].

Analyzing the delay within a domain will therefore depend on the route selected, which can be expressed as, the summation of delays through the components in the route [10]. The link delays are easily calculated since they are proportional to the bandwidth. However the equipment delays are more difficult to quantify. A router is basically a specialized computer system with additional complexity introduced due to the real-time operation of the router OS [11].

1.1. Quantifying delays within a router

A simple laboratory network was set up with the use of a dual ported Linux machine running Wireshark as a method of measuring delays across a router. An initial experiment was conducted to identify the accuracy of the measurements and this delay which should be 0 μsecs, was on average 9 μsecs. This would be the error bar for the network.

Packets which enter a router via the network interface card are filtered by their destination network address using the routing table. The delay of this process is dependent on the software to setup the process and the hardware components e.g. memory access time. Performance of router hardware is highly variable since it is dependent on the underlying technology, including the processing power and memory capacity. Additionally, high throughput hardware can be purchased which exhibits performance improvements due to the specification. Networks typically contain equipment of varying ages which results in performance variations. In this work, to enable other factors to be compared, consistent typical performance hardware has been used.

Router OS are optimized for routing of packets however they are also required to perform many other tasks which will be dependent on the feature set. A comparison of OS size and number of supported/ running processes was undertaken using an OS with basic functionality and another with advanced functionality. Applying a configuration that only enabled the interfaces and OSPF routing protocol to the router, the show processes command provided the information in Table 1. Clearly the advanced OS runs many more processes for a given configuration compared to a basic OS which will have an effect on the responsiveness of the CPU and the amount of memory required.

Table 1

OS Comparison

Functionality	OS Size	Number of Processes	Active Processes > 2
Basic	12MBytes	73	32
Advanced	29MBytes	184	51

If a core part of the OS is enhanced with additional functionality e.g. HTTP or DHCP Servers it can have an additional adverse effect on the size of the OS and its performance to that seen in Table1.

1.2. Measurement of delays within a router

Identical tests were undertaken using ICMP packets, to quantify the delay across a router with a basic OS and then with an advanced OS. The variation in times in the individual results, is due to the number of processes in the OS. Results were analyzed using histogram techniques and plotted in Fig. 2 where the x-axis shows the measured delay in μsecs. and the y axis shows the number of times this value was obtained.

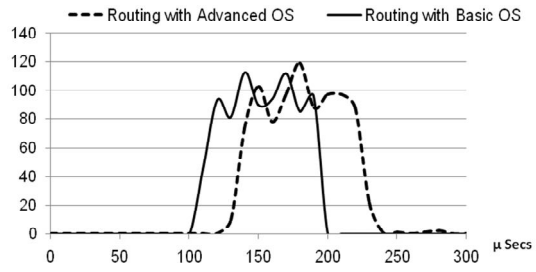

Fig. 2. Delay through router with different OS

Security is typically implemented on a router using ACLs. For an individual packet each rule is evaluated in turn until a matching rule is found. Standard ACLs only filter on the source IP address of a packet whereas extended ACLs provide the capability to filter on additional fields such as destination address, protocol and port numbers. There are many good texts available on the subject of ACLs so it is not covered in this paper [12].

Measurements were taken to investigate the delays when the router, running the basic OS, was configured with 100 rule ACLs (Fig. 3).

As expected no ACL gave the least delay, a standard ACL increased the delay by approximately 110% and the extended ACL an increase of 270%.

Further work was carried out to investigate the delay experienced by packets matched against an increasing number of rules. Fig. 4 shows for a Basic OS increasing the number of rules in the list has a significant effect on the delay. Increasing the number of rules in an Extended ACL from 100 to 1000 has the effect of increasing the delay by approximately 530%.

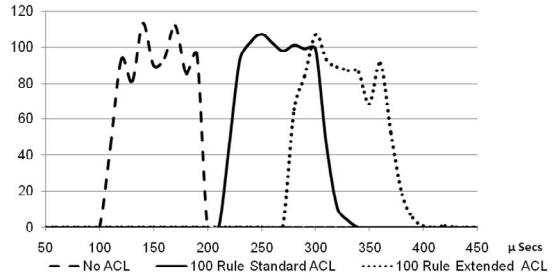

Fig. 3. Delay through router running Basic OS

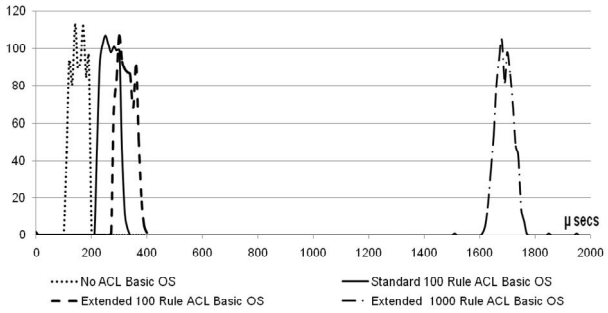

Fig. 4. Delay through router with Basic OS

Repeating the experiment using the same 100 and then 1000 rules did not incur any additional delay using an Advanced OS (Fig. 5). This performance improvement could not be due to hardware so it must be due to a software enhancement in the router OS. Cisco do not release details of the OS however it is likely that a binary decision technique has been employed because the delay time is not dependent on the number of rules that are in the ACL.

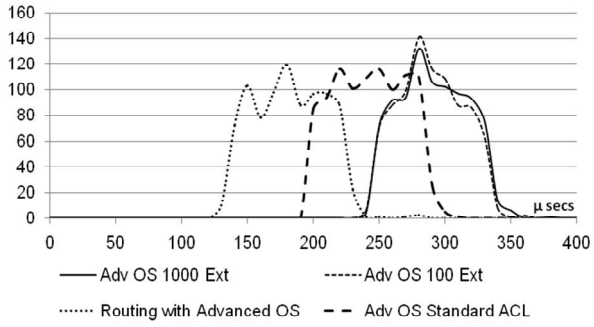

Fig. 5. Delay through router with Advanced OS

2. Analysis of delays within a Router

An average value of the range was calculated in order to provide a single value associated with each test. Results show that some parameters have a greater significance than others (Table 2).

By using a router with a basic OS rather than an advanced OS it can be seen that standard routing is faster by around 15%. Configuring extra services on the advanced OS is expected to further increase the latency. When ACLs are configured, for a basic OS the average delay is increased by around 80% for a standard ACL and 110% for an extended ACL.

Table 2

Average delays for all tests (times in µs)

IOS version	No ACL	Standard	Ext 100	Ext 1000
Basic	150	271	320	1685
Advanced	172	239	300	309

When using a router with a basic OS adding more rules to an ACL has a significant effect on the delays which can be of the order of 1400% for 1000 rules. The advantage of the advanced OS functionality is that the number of rules used in an ACL does not have an effect on the delay (Fig. 6).

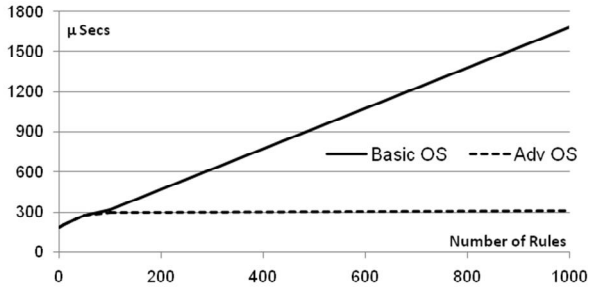

Fig. 6. Delay v number of rules

3. Delays within a Domain

Within a domain either static routes are configured or a routing protocol is used to select a route. Theoretically, the cumulative delay (D_d) for a given path can be calculated by the summation of the delays in the equation for each router (n) in the route.

$$D_d = \sum_{i=1}^{n} D_{h_i} + \sum_{i=1}^{n} D_{os_i} + \sum_{i=1}^{n} D_{a_i} + \sum_{i=1}^{n} D_{s_i} +$$

$$+ \sum_{i=1}^{n} D_{ta_i} + \sum_{i=1}^{n} D_{nr_i} + \sum_{i=1}^{n} D_{p_i} + \sum_{i=1}^{n} D_{q_i} + \sum_{i=1}^{n} D_{r_i},$$

D_h – router hardware, D_{os} – router OS, D_a – applications, D_s – services, D_{ta} – type of ACL, D_{nr} – delay per ACL rule, D_p – protocol, D_q – queuing delay variation, D_r – total delay.

3.1. Calculation for Example route

For example an optimized route selected from Fig. 1 between the Wireless connections and the internet contains 7 nodes and can be used to investigate the implications of the values in Table 2.

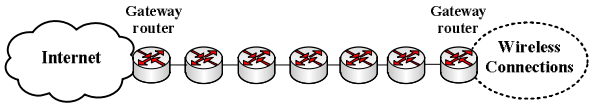

Fig. 7. Simplified Route through Domain

For a typical security policy, the domain gateway routers (ingress & egress) may have 1000 rules configured. In addition, each internal router may have 100 rules configured, possibly on both ingress and egress interfaces within each router.

Fig. 7 shows a simplified version of the route and Table 3 shows the expected delay for a packet using both a basic and advanced version of the OS. Analyzing the delay within a domain will depend on the route selected, which can be expressed as the summation of delays through the components in the route. The worse case is all routers configured with ACLs and the basic OS utilized.

By simply using an advanced OS in all the routers, an improvement of 135% can be obtained over a basic OS. Optimizing the route based on the conditions derived from Tables 2, the best case is using advanced OS everywhere an ACL is applied and a basic ACL where no ACLs are applied. If this principle is used in a trusted domain then routers loaded with the advanced OS would be used at the ingress/egress points and routers with basic OS and no ACLs within the rest of domain. The best case shows an improvement of 260% over using a basic OS everywhere in the domain and 50% over using an advanced OS everywhere.

Table 3

Calculated Delays for example network

IOS version	No ACL	Ext 100	Ext 1000	Total μsecs
Basic OS	0	1600	3370	4970
Basic Optimized	750	0	3370	4120
Advanced OS	0	1500	618	2118
Advanced Optimized	860	0	618	1478
Mixed Optimized	750	0	618	1368

3.2. Condition controlling optimization

If all the traffic within a domain is trusted then it should be possible to eliminate the ACLs from routers within the Domain on condition that the security is catered for in the ingress and egress points of the domain. It is imperative that any optimization process that is undertaken preserves the security policy for the domain. Therefore it is required that all ingress/egress points in the domain are capable of identifying and removing all packets which are not permitted by the security policy. If the gateway routers are configured with an advanced OS the number of rules generated is insignificant.

4. Protocols and the placing of ACLs

Most routers provide the ability to configure ACLs using different routed protocols. In addition to IP, ACLs can also be configured for IPv6, IPX, AppleTalk etc and each protocol is configurable on each interface. Usually routers only allow a single ACL to be configured on a router interface for a given direction and protocol. IP in its IPv4 guise is the predominant protocol used in network communications and on the Internet, therefore this is the only protocol considered in this study.

It is not a simple process to replace all the ACLs in routers internal to the domain with rules at the ingress and egress points. An investigation on the feasibility of carrying out this process has been considered for the simplest case of a standard ACL. There is a possibility that anomalies such as redundancies may exist within an ACL which could be removed without affecting the functionality of the ACL. This principle can be extended to consider subsequent routers along a path.

The basis for optimization is that if some addresses are denied access by subsequent routers then these can be moved to routers earlier in the route. Additionally the range of rules can be extended to be incorporated into earlier rules by modifying the mask. The effect would be the reduction of the overall delay across the domain. This has to be done with great consideration to ensure that the security policy is not violated. In a simple example with 3 routers the effect would be a reduction in the delay over the domain by around 65%.

4.1. Processing of Rules

To prove the possibility of this being applied in practice the processing of the rules has been investigated. ACLs provide a very simple decision process since for every packet tested there are only 2 possibilities, permit or deny which are defined by the first rule match.

When considering all the rules in an ACL it is possible to represent the result for all possible IPv4 packets by constructing an array which contains a bit representing each possible IP address. This was done by developing a custom-built data structure. The contents of the array is initially set to reflect all addresses being denied which corresponds to the implicit “deny any any” statement found at the end of all ACLs.

The process starts by using the ACL found at the ingress point of the domain, the array is populated with the filtering action for each possible IPv4 packet. The rules are evaluated starting from the bottom of the list since the priority of a rule increases for rules higher in the list. For each router containing an ACL along the network path the array is rewritten based on the rules defined in each ACL. When the first rule in the ACL has been reached then the array reflects all the rules of the ACL. Then the array is modified to reflect the rules found in the ACLs of the next router in the path. This process is continued until the final point in the domain is reached. The final content of the array represents the security policy for that route through the domain.

An intermediate step can be performed to merge rules into larger ranges providing they have the same filtering action. Each time a change is made in the security policy for a particular path through the domain it would be necessary to repeat the above steps to obtain a new ACL which reflects the contents of the rewritten array.

4.2. Creation of new rules

Having created the array it is necessary to do a reverse transformation to produce the rules for an ACL which replicates the functionality of the distributed ACLs for a given filtering direction. The resulting ACL can be further optimized to remove any redundancies which may be present. It may also be possible to consolidate similar rules using wildcard masks and range commands for port numbers. Once fully optimized, the ACL can be applied to the ingress/egress. The final list may be considerably longer than the initial ACL but based on the work carried out above by implementing this in a router with an advanced OS the additional delay is insignificant.

One of the main concerns that network administrators have about this technique is that the final list can bear little relationship to the original ACL created. To alleviate this concern the original ACL is kept and a cross reference list is provided showing the correlation between the rules in the new ACLs.

Conclusion

Utilizing routers within a domain to provide security does have an impact on the performance of the network since it introduces significant delays due to the equipment. There are some relatively simple steps that can be taken to improve the performance.

By investigating the theoretical aspects of delays through routers and carrying out a series of measurements it has been possible to improve the model of delays encountered by a packet as it transverses a domain. It has also been possible to quantify the delays to understand which components are more significant which leads to a series of rules that can be used as best practice when designing large networks.

There are significant variation in the delays experienced using different versions of the OS in the router. A more advanced OS adds delays to the basic routing process but if other functionality is required then an advanced OS has to be used.

Optimal performance can be gained by not having ACLs enabled in a router. It is not possible to remove the ACLs from all routers within a domain but there are gains to be made by reducing the number of routers that have ACLs enabled. By using an Advanced OS the number of rules in an ACL is insignificant. Since a domain has a common security policy then it should be possible to optimize the placement of ACL rules to ensure that the minimum number of routers in a domain use an ACL.

Having completed optimization on the number of routers requiring an ACL then using a basic OS for a router without ACLs and using advanced OS for the routers that do require an ACL will show an overall improvement of performance.

By investigating a domain which has a single security policy and therefore allows functionality of routers to be allocated this paper considers how it can be optimized. It shows that by ensuring the OS with the appropriate functionality is used an improvement of performance will be gained. By utilizing an OS with advanced functionality an improvement around 130% is possible.

Furthermore this paper shows that by moving the ACLs to only the egress/ingress points of a domain that a performance improvement of the order of 250% can be gained over using a basic OS or in excess of 50% over using an advanced OS.

The area that would be considered for future work include: the effect of more advanced hardware, effect of using additional functionality / services to the network within a router. More importantly the use of IPv6 in the internet is a far more complex issue due to the size of the parameters involved and so requires special investigation.

References

1. El-Atawy, A. Adaptive Statistical Optimization Techniques for Firewall Packet Filtering [Text] / A. El-Atawy, H. Hamed, E. Al-Shaer // IEEE Infocom. – 2006. – Barcelona, Spain, 23-29 April 2006. – P.1 – 12

2. Davies, J.N. Improving the Performance of IP Filtering using a Hybrid Approach to ACLs [Text] / J.N. Davies, V. Grout, R. Picking // Proceedings of the 8th International Network Conference (INC2010) Heidelberg, Germany, 6 - 8 July 2010. – P. 32 – 41/

3. Rule Dependencies in Access Control Lists [Text] / V. Grout, J. McGinn, J.N. Davies, R. Picking, S. Cunningham // Proceedings of International Conference WWW/Internet (IADIS) San Sebastian, Spain 26-28 February 2006. – P. 120 – 132/

4. Al-Shaer, E.S. Modelling and Management of Firewall Policies [Text] / E.S. Al-Shaer, H.H. Hamed // IEEE Transactions on Network and Service Management. – April 2004. – Vol. 1, no.1. – P. 2 – 10.

5. Hari, B. Detecting and Resolving Packet Filter Conflicts [Text] / B. Hari, S. Suri, G. Parulkar // Proceedings of the 19th Joint Conference of the IEEE Computer and Communications Societies (INFOCOM00), Tel Aviv, Israel, March 26-30, 2000, IEEE. – Vol. 3. – P. 1203 –1212.

6. Analysis of point-to-point packet delay in an operational network [Text] / B.Y. Choi, S. Moon, Z. Zhang, K. Papagiannaki, C. Diot // Computer Networks. – September 2007. – № 51. – P. 3812 – 3827.

7. Moy, J. RFC 2328 OSPF Version 2 [Text] / J. Moy // The Internet Society. OSPFv2. – 1998. – 120 p.

8. Resende, M.G.C. Handbook of Optimization in Telecommunications [Text] / M.G.C. Resende, P.M. Pardalos. – Springer Science + Business Media, New York, NY, 2006. – 320 p.

9. Hohn, N. Capturing router congestion and delay [Text] / N. Hohn, K. Papagiannaki, D. Veitch // IEEE/ACM Transactions on Networking 17. – 2009. – Vol. 17, No.3. – P. 789-802.

10. Lai, K. Measuring link bandwidths using a deterministic model of packet delay [Text] / K. Lai, M. Baker // Proceedings of the conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM '00) Stockholm, Sweden, August 28-September 1 2000. – P. 45 – 52.

11. Bollapragada, V. CCIE Professional Development: Inside Cisco IOS Software Architecture [Text]/ V. Bollapragada, R. White, C. Murphy. – Cisco Press, 2000. – 13 p.

12. Sedayao, J. Cisco IOS Access Lists [Text] / J. Sedayao // O'Reilly & Associates, Inc., Sebastopol, CA, USA 2001. – P. 22 – 31.

Поступила в редакцию 23.02.2012

Рецензент: д-р техн. наук, проф., зав. каф. информационных управляющих систем О.Е. Федорович, Национальный аэрокосмический университет им. Н.Е. Жуковского «ХАИ», Харьков, Украина

ПРИНЦИПИ ОПТИМІЗАЦІЇ СПИСКІВ КОНТРОЛЮ ДОСТУПУ В МЕЖАХ ДОМЕНУ

Д.Н. Девіс, П. Комерфорд, В. Граут, Н. Рвачова, О. Корх

Інфраструктура великих мереж характеризується розбиттям на області, що мають спільну політику безпеки та називаються доменами. Політика безпеки в домені, зазвичай, реалізується в усіх вузлах, що призводить до затримок, пов'язаних із фільтрацією пакетів та негативно впливає на продуктивність мережі. Застосування списків контролю доступу (ACL) в маршрутизаторах призводить до збільшення накладних витрат. Загальна продуктивність мережі може бути підвищена, якщо перевірка пакетів здійснюватиметься лише в граничних точках домену, виключаючи перевірку ACL всередині домену, при цьому точки входу/виходу повинні мати еквівалентну функціональність.

Ключові слова: домен, продуктивність, затримка в маршрутизаторі, список контролю доступу, ACL оптимізація, фільтрація IP-пакетів.

ПРИНЦИПЫ ОПТИМИЗАЦИИ СПИСКОВ КОНТРОЛЯ ДОСТУПА В ПРЕДЕЛАХ ДОМЕНА

Д.Н. Девис, П. Комерфорд, В. Граут, Н. Рвачева, О. Корх

Инфраструктура больших сетей характеризуется разбиением на области, которые имеют общую политику безопасности и называются доменами. Политика безопасности в домене, как правило, реализуется во всех узлах, что приводит к задержкам, связанным с фильтрацией пакетов и негативно влияет на производительность сети. Применение списков контроля доступа (ACL) в маршрутизаторах приводит к увеличению накладных расходов. Общая производительность сети может быть повышена, если осуществлять проверку пакетов только в приграничных точках домена, исключая проверку ACL в середине домена, при этом точки входа/выхода должны иметь эквивалентную функциональность.

Ключевые слова: домен, производительность, задержка в маршрутизаторе, список контроля доступа, ACL оптимизация, фильтрация IP-пакетов.

Девис Джон Н. – д-р, университет Глиндор, Рексем, Великобритания, e-mail: j.n.davies@glyndwr.ac.uk.

Комерфорд Пол – аспирант, Glyndwr University, e-mail: p.comerford@glyndwr.ac.uk.

Граут Вик – д-р техн. наук, профессор университет Глиндор, Рексем, Великобритания, e-mail: v.grout@glyndwr.ac.uk.

Рвачева Наталия – канд. техн. наук, старший преподаватель кафедры компьютерной инженерии Полтавского национального технического университета имени Юрия Кондратюка, Полтава, Украина, e-mail: rvacheva_n@mail.ru.

Корх Олег – канд. техн. наук, старший преподаватель кафедры прикладной математики, информатики и математического моделирования Полтавского национального технического университета имени Юрия Кондратюка, Полтава, Украина, e-mail: korkholeh@gmail.com.