Evaluation of Tailored Mutation Operator in a Parallel Genetic Algorithm for Pavement Maintenance Treatment Scheduling

Glen Cancian, and Wayne Pullan

***Abstract*— The maintenance of an existing large road network is a key focus area for road authorities around the world. The pressures associated with the ever-increasing road network and often shrinking budgets means that it is essential that road authorities invest maintenance budgets wisely. In line with this objective, most road authorities' employee a Pavement Management System (PMS) to assist in making maintenance decisions. PMSs must solve a very large optimization problem involving thousands of road segments with multiple possible treatments. There is a wide range in the cost of these treatments and also in the magnitude and duration of their improvement. The optimization problem is to identify a minimum cost, 20-year maintenance program that ensures all segments are maintained at an acceptable level (which varies depending on factors such as the amount of traffic and the type of traffic). In addition to the 20-year overall budget, there are yearly budgets constraints which must be met and many other constraints such as the availability of staff and machinery.**

Previous research has shown significant benefit arises from the adoption of a genetic algorithm-based PMS. This paper builds on this research through the application and evaluation of a tailored, parallel genetic algorithm within a PMS. A tailored genetic algorithm is evaluated using a real-world road network of 1,335 road segments executed using 12 processing units with annual budgets ranging between $40 and $50 million. Over a total of 174 trials, the tailored genetic algorithm was 46% more successful than a standard genetic algorithm at producing an optimised program of works that satisfied all budget constraints, typically with a lower overspend.

***Index Terms*— pavement management system, parallel genetic algorithm, tailored mutation, budget constrained genetic algorithm**

I. Introduction

ROAD authorities are responsible for programming maintenance and rehabilitation treatments for road networks to best utilise available funding. These authorities use a Pavement Management System (PMS) as a decision support tool to assist with identifying optimal treatment programs over multi-year planning horizons. Historically PMSs have been labour intensive manual methods however many road authorities now find it advantageous to employ a computerised PMS to take advantage of their significantly improved processing power. These computerised PMS are comprised of six modules [1,2,3] at the core of which is an optimisation module. The optimisation module processes input data to produce a schedule of programmed maintenance and rehabilitation treatments for multiple road segments over a multi-year period. This input data is typically either collected utilising network survey vehicles on an annual or bi-annual basis or is a function of the asset over time (e.g. age). The attributes collected normally consists of ride quality, strength and deterioration indicators like roughness, rutting and cracking. This data is used by road authorities to prepare forward works programs for design and subsequent delivery in future years.

Over time a variety of approaches have been implemented in the PMS optimisation module, a significant number of these have been prioritisation techniques rather than exhibiting optimisation attributes. The approaches that have been employed in the optimisation module can be categorised into two distinct groups, deterministic and stochastic approaches (the need for stochastic optimization arises because of the extremely large solution space which is of order 20 * 81335 for the road network used in this paper).

Those PMS currently employed by industry are more likely to be deterministic. These include decision trees [4,5,6] and their variants [7], expert systems [8,9,10,11], knowledge-based systems [12], linear and dynamic programming [3,13,14,15]. Stochastic techniques have mainly been the subject of research and development, including Markov decision models [16], particle swarm optimisation [17] and genetic algorithms [18,19].

Genetic algorithms have been shown to produce an optimised schedule of programmed maintenance and rehabilitation treatments for a real-world pavement network [20]. The optimised schedule significantly reduced the value of investment required to maintain the road network over the analysis period.

This study builds upon the work undertaken by Cancian, Chai and Pullan [21] whose Parallel Genetic Algorithm (PGA) based pavement management treatment scheduling system employed a standard, randomised mutation operator. Through a number of computational experiments this study evaluates a PGA with a tailored mutation operator designed to mutate parts of solutions that contribute to noncompliance with the objective function. The results

Manuscript received June 15, 2020; revised August 31, 2021.

G. Cancian is the Manager of Asset Management System & Strategy for DM Roads, Milton, QLD 4064 Australia phone: +61409 607 072; e-mail: glen.cancian@dmroads.com.au.

W. Pullan is with the School of Information and Communication Technology at Griffith University, Southport, QLD 4222 Australia (e-mail: w.pullan@griffith.edu.au).

Solution Vector:

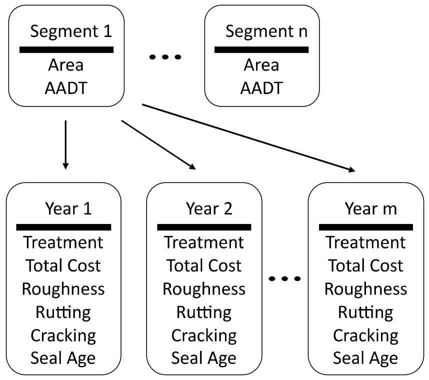

Tailored Mutation:

```
if (sum of solution overspend > 0) then
    for (y = 0 to number of years)
        if (budget need > annual budget) then
            Find segment treated with a high cost
              treatment
            Replace treatment with lower cost treatment
            Process remainder of segment using rule
              base
            exit subroutine
        end if
    end for
end if
```

Fig. 1. Solution vector that maintains the state of each road segment, the treatment to be applied and the total treatment cost. Also shown is the pseudo-code for the tailored mutation operator which operates on each solution vector to replace the higher cost segment treatments with lower cost options.

generated by the tailored PGA using the same real-world road network of 1,335 road segments are outlined and analysed; and the effectiveness of the tailored mutation is evaluated.

The remainder of this paper is organised as follows: in the following section the tailored PGA treatment scheduling technique is presented, then experimental results generated using incremented budget constraints are reported and analysed. A comparison is made and the effectiveness of the tailored mutation is highlighted. Then finally a summary is presented along with opportunities for future research in the conclusion.

II. Tailored Parallel Genetic Algorithm

As is the case in the original study by Cancian, Chai and Pullan [21], the tailored PGA treatment selection technique presented in this study closely follows a standard genetic algorithm, however it is implemented across multiple processors [22].

In a standard genetic algorithm a pool of solutions is maintained and manipulated through a series of crossover and mutation operations over a specified duration or until a termination criteria is met. These operations facilitate the improvement of solutions in the pool using a specified objective function. A PGA employs the same operations however the solution pool is managed by a Master process while the crossovers and mutations are performed by Slave processes.

The tailored PGA algorithm presented in the study closely aligns with the algorithm in the original study with the exception of a tailored mutation operator. Fig. 1. includes a representation of the solution vector and details the algorithm for the tailored mutation operator. This modification takes the form of a subroutine executed by each Slave process. The tailored mutation initially identifies the first year in which the annual budget is exceeded. Then selects a segment that is programmed to be treated with one of the three highest cost treatments (rehabilitation, correct and 45mm overlay or 45mm overlay) in that year. Once selected, this high cost treatment is replaced with one of the remaining lower cost treatment alternatives and the remainder of the analysis period processed in accordance with the rule-base rules.

III. Computational Experiments

The objective of these computational experiments was to compare the effects of the tailored mutation operator with those of the original study which employed a standard GA with a randomised mutation operator. These effects will be evaluated using the following criteria:

A: The percentage of viable solutions generated that satisfied all maximum condition and annual budget constraints;
B: Spend profiles of viable solutions;
C: Propagation of viable solutions through the solution pool;
D: New solution generation;
E: Solution composition; and
F: Effectiveness of mutation operator on budget overspend objective.
G: Segment treatment plan generation

For comparability of results, the computational experiments were undertaken using the same road network as the original study comprising of 1,335 segments. Each segment is surfaced with either an asphalt or sprayed seal surface and recursive deterioration models are used to deteriorate segment condition in those years that a treatment isn't triggered taking account of the segment's environment, material composition and average annual daily traffic volume.

After the application of the mutation operator a 29-rule rule-base was used to process the remainder of each segment. Within this rule-base there are multiple rules that trigger a single Segment Treatment (ST) with eight individual treatments being used in this analysis, ST1 - ST8. The unit rates for each treatment vary between \$58.00/m2 for ST1 and \$6.00/m2 for ST8 and the condition attributes used in the analysis are roughness, rutting, cracking and seal age.

The solution pool used in the tailored PGA contained 10 solutions and admission criteria was two-fold. If the total overspend against the annual budget was greater than zero for any solution in the solution pool, the new candidate solution was added if it had a lower total overspend. Alternatively, if each solution in the pool satisfies the annual budget constraint then a candidate solution would be admitted to the pool if it had a lower total spend than the

ISBN: 978-988-14049-1-6
ISSN: 2078-0958 (Print); ISSN: 2078-0966 (Online)

worst solution in the pool.

Maximal condition constraints were applied to the roughness, rutting and seal age attributes. Due to the poor initial condition of the road network, each maximum condition constraint was checked from year two of the analysis period. The maximum condition constraints were hard constraints. Each of these computational experiments were executed over an analysis period of 20 years.

For each trial, the tailored PGA was executed for 30-minutes with an annual budget constraint of either $40 million, $45 million or $50 million using 12 processors. These 3 individual combinations were executed with 58 seed values (executed twice with 29 unique seeds), resulting in a total of 174 individual tailored PGA trials. These budget values were selected as the PGA trials from the original study with budgets between $25 million and $35 million failed to produce solutions that satisfied all constraints. The seed values took the form of large prime numbers with each Slave seeded with large prime plus world rank where world rank is the Slave’s rank among its peers.

When comparing each evaluation criteria in the sub-headings below, only those trials from the original study executed using 12 processors are considered.

A. *Viable Solution Generation*

The tailored PGA provided at least one viable solution that satisfied all maximum condition and annual budget constraints in 27 trials of the total 174 undertaken. In comparison, considering only those trials with an annual budget constraint of $40 million, $45 million or $50 million and executed using 12 processors from the original study, 7 trials produced at least one viable solution that satisfied all constraints from 66 trials undertaken.

Taking account of the difference in number of trials, the study with the tailored mutation operator produced 46% percent more trials with at least one viable solution that satisfied all maximum condition and annual budget constraints.

TABLE 1
SOLUTION PROPAGATION

Propagation	Standard	Tailored
2	2	-
3	-	-
4	-	-
5	-	1
6	-	-
7	-	-
8	-	6
9	4	5
10	1	15
Total	7	27

Viable solution propagation through solution pool. Clearly the tailored mutation consistently produced more viable solutions whereas the standard mutation was more erratic.

All of the successful trials using the tailored mutation operator that produced solutions satisfied the $50 million annual budget constraint using processors. Similarly, only one of the successful trials from the original study produced a solution that satisfied the $45 million budgetary constraint, all others satisfied an annual budget of $50 million.

B. *Spend Profiles of Viable Solution*

Fig. 2. panels a and b depict the annual spend profile of each of the trials that generated viable solutions from the original and current studies. It is evident that although executed with a different seed value, each viable solution generally follows a similar annual spend profile, irrespective of the mutation operator employed in each trial.

This pattern is likely heavily influenced by the present

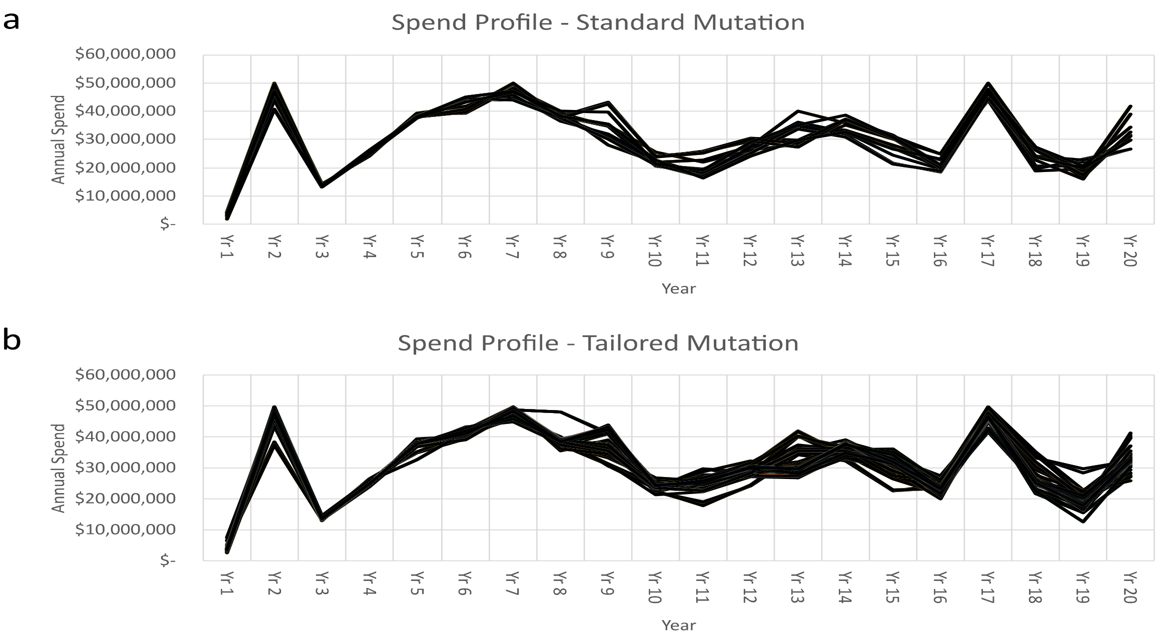

Fig. 2. Year by year spend profiles for the 27 viable solutions. All viable solutions started with an initial high spend to fix immediate issues and then followed the same general pattern of spending.

TABLE 2
NEW SOLUTIONS

New Solutions	Standard	Tailored
100 – 150	-	-
151 – 200	1	-
201 – 250	1	5
251 – 300	1	10
301 – 350	2	9
351 – 400	1	2
401 – 500	1	1
Total	7	27

Number of new solutions included in the solution pool for each trial that produced viable solutions showing that the tailored mutation was consistently able to produce a higher number of viable solutions at each generation..

condition of the road network; the deterioration models applied to each pavement segment in years when no treatment has been triggered; and the treatment resets applied after each treatment.

C. Propagation of Solutions

Table 1 details the count of pool solutions within each trial group that adhere to all maximum condition and budget constraints.

Evident from the data is that those trials executed using the tailored mutation had excellent solution propagation. Only one of these trials failed to attain at least eight viable solutions that satisfied all constraints within the solution pool.

This, however, contrasts with the seven trials from the original study that provided viable solutions. This outlines that those from the original study exhibited a reduced level of propagation in comparison to the current study with the tailored mutation.

D. New Solution Generation

Table 2 details the number of new solutions added to the pool by each of the trials that generated viable solutions in both studies. All but a single trial from the original study produced greater than 200 new solutions that were added to the solution pool.

TABLE 3
RETURNED SOLUTIONS (%)

% Returned	Standard	Tailored
40 – 45	-	1
45 – 50	-	1
50 – 55	4	5
55 – 60	-	4
60 – 65	2	7
65 – 70	1	7
70 – 75	-	2
Total	7	27

Percentage of solutions returned by Slave processors included in the solution pool for each trial that generated viable solutions.

TABLE 4
SOLUTION COMPOSITION

Treatment	Standard	Tailored	Treatment Category
ST1	848	868	
ST4	3	3	Higher Cost
ST6	269	255	
ST2	102	86	
ST3	2	1	
ST5	218	211	Lower Cost
ST7	66	63	
ST8	201	181	
Total	1710	1668	

ST1 = rehabilitation, ST2 = correct and seal, ST3 = correct and 30mm overlay; ST4 = correct and 45mm overlay, ST5 = 30mm overlay, ST6 = 45mm overlay, ST7 = fabric spray seal, ST8 = reseal

Table 3 lists the percentage of solutions returned by the Slave processors that were actually added to the solution pool by the Master. Except for a two trials, all trials that returned viable solutions have a return rate of between 50 and 75%. Those returned solutions that weren't added to the solution pool were discarded by the Master processor during the course of its improvement and uniqueness verifications.

This was because, in the time between the Slave checking the new solution is an improvement on the recorded value of the current worst solution in the pool, the pool had already been updated with solutions produced by other Slaves.

E. Solution Composition

The average number and type of each treatment included in each trial that produced viable solutions are tabulated in Table 4.

Evident from the table is that those trials with the tailored mutation operator triggered five more higher cost treatments and 47 fewer lower cost treatments than those using the standard mutation. An overall reduction of 42 treatments over the 20-year analysis period.

F. Effectiveness of Mutation

The effectiveness of the tailored mutation operator can be evaluated by analysing the difference in total solution overspend immediately before and after a mutation operation. A 30-minute trial employing 12 processors was undertaken and the results analysed.

Due to the additional computational expense required to perform each tailored mutation, only 177 mutations were performed in the analysis period while 367 standard mutations were undertaken. The average difference in overspend following the randomised mutation utilised by Cancian, Chai and Pullan [21] was \$3.547 million with a standard deviation of $118 * 10^6$. In comparison following the tailored mutation operation the average difference in overspend was -\$1.697 million with a standard deviation of $103 * 10^6$.

ISBN: 978-988-14049-1-6
ISSN: 2078-0958 (Print); ISSN: 2078-0966 (Online)

TABLE 5
ALTERNATE TREATMENT PLANS

Treatment Plans	Standard	Tailored
1	550	283
2	573	595
3	87	214
4	119	96
5	6	91
6	-	18
7	-	35
8	-	2
9	-	1
Total	1335	1335

Number of alternate treatment plans showing the wider range of treatment plans utilised when the tailored mutation operator is used.

Despite producing less solutions during the 30-minute trial than the standard mutation operator, the tailored mutation operator produced solutions that offer greater improvement. The most improved solution produced following a standard mutation operator has a total overspend of $37,208,982, in comparison the tailored mutation operator produced an improved solution with an overspend of $9,242,000.

G. Segment Treatment Plans

Table 5 details the number of alternate treatment plans generated for each of the 1,335 network segments in both studies. For the original study, 84% of segments only had 1 or 2 treatment plan alternatives in those trials that generated viable solutions. In contrast, only 66% of the trials that employed the tailored mutation operator had less than 3 treatment plan alternatives with all other segments having between 3 and 9 alternate treatment plans.

The tailored mutation operator ensures the PGA produces a greater variety of alternate treatment plans that are then ultimately aggregated into whole of network solutions that satisfy budget and condition constraints.

IV. CONCLUSION

This paper presented a budget constrained tailored implementation of a Parallel Genetic Algorithm (PGA) based pavement management treatment scheduling system along with results from computational experiments generated using a real-world road network of 1,335 road segments. The computational experiments were undertaken using 12 processors over a series of 30-minute trials and were compared with the results of a previous study that matched the specifics of this study however employed a randomised mutation operator.

Like the previous study, the resulting PGA with a tailored mutation operator produced an optimised pavement maintenance and rehabilitation program of works that satisfied maximum condition constraints placed on roughness, rutting and cracking, in addition to satisfying an annual budget constraint of $50 million when executed for 30-minutes.

The computational experiments have highlighted that the tailored PGA algorithm provided a 46% increase in the number of trials that generated viable solutions; in addition to providing better solution propagation of viable solutions through solutions pools.

There was no discernible difference between both studies in the number of new solutions added to each solution pool in the trials nor the percentage of returned candidate solutions that were actually added to the solution pool.

In addition to the above, despite producing less candidate solutions in the same trial period due to the additional computational expense of performing the tailored mutation, those candidates from the tailored mutation exhibit a greater reduction in overspend pre vs. post mutation, and less variance from the mean (lower standard deviation).

Opportunities for further research include increasing the number of trials, or alternative maximum condition constraints could be implemented and an alternative condition based objective function could be optimised. Both the original study and this study have used a rule base to process solutions following the application of a treatment. Variations in this rule base or alternatively not using any rules could lead to different treatment profiles and are worth further investigation.

REFERENCES

[1] AASHTO and Task Force on Pavements, Pavement Management Guide - Executive Summary Report, American Association of State Highway and Transportation Officials, Washington DC, USA, 2001
[2] Centre for Pavement Engineering Education (CPEE), CPE 607 - Pavement Management Reading, 1:13, 2009
[3] H. L. Xiong, Miao and Q. Shi, Network level framework for Beijing Roadway Management System, In Proc. IEEE Conference on Intelligent Transportation Systems, 1309-1312, vol. 2, Shanghai, China, October 12-15, 2003
[4] R.G. Hicks, S.B. Seeds, and D.G. Peshkin, Selecting a Preventive Maintenance Treatment for Flexible Pavements, US Department of Transportation Federal Highway Administration Office of Asset Management, Washington DC, USA, August, 2000
[5] D.G. Peshkin and T.E. Hoerner, Pavement Preservation: Practices, Initiatives and Research Plans, National Cooperative Highway Research Program, Project 20-07, Washington, USA, May, 2005
[6] G. Zhou and L. Wang, Co-location decision tree for enhancing decision-making of pavement maintenance and rehabilitation, Transportation Research Part C: Emerging Technologies, Elsevier Ltd., vol. 21, Issue 1, April, 2012, pp. 287-305
[7] G. Cancian, G. Chai and W. Pullan, Pavement Maintenance Treatment Selection Using a Tree-Based Selection Algorithm, TRB 94th Annual Meeting Compendium of Papers, Transport Research Board, Washington DC, United States, January, 2015
[8] J.G. Cabrera and M. Al-Shawi, Performance of concrete pavements using "pavement expert", In Proc. Conference on the Application of Artificial Intelligence Techniques to Civil and Structural Engineering, Civil-Comp Press, Edinburgh, UK, 1987, pp 131-135
[9] J.J. Hajek, G.J. Chong, R.C.G. Haas and W.A. Phang, Rose: A knowledge-based expert system for routing and sealing, In Proc. 2nd North American Conference on Managing Pavements, Ontario Ministry of Transportation, Canada, 1987, pp. 2.301-2.313
[10] W.V. Harper and K. Majidzadeh, Use of expert opinion in two pavement management systems, Transportation Research Record: Journal of the Transportation Research Board, Washington DC, USA, No. 1311, 1991, pp. 242-247
[11] S.G. Ritchie, C.I. Yeh, J.P. Mahoney and N.C. Jackson, Surface condition expert system for pavement rehabilitation planning, Journal of Transportation Engineering, American Society of Civil Engineers, Virginia, USA, vol. 113, issue 2, March, 1987, pp. 155-167
[12] N. Ismail, A. Ismail and R. Atiq, An overview of expert systems in pavement management, European Journal of Scientific Research, Scientific Research Platform, vol. 30, May, 2009, pp. 99-111

[13] K.J. Feighan, M.Y. Shahin and K.C. Sinha, A dynamic programming approach to optimization for pavement management systems, In Proc. 2nd North American Conference on Managing Pavements, Ontario Ministry of Transportation, Canada, 1987, pp. 2.195-2.206

[14] K.J. Feighan, M.Y. Shahin, K.C. Sinha and T.D. White, Application of dynamic programming and other mathematical techniques to pavement management systems, Transportation Research Record: Journal of the Transportation Research Board, Washington DC, USA, No. 1311, 1988, pp. 90-98

[15] T.F. Fwa, K.C. Sinha and J.D.N. Riverson, Highway routine maintenance programming at network level, Journal of Transportation Engineering, American Society of Civil Engineers, Virginia, USA, vol. 114, issue 5, September, 1988, pp. 539-554

[16] F. Guignier and S. Madanat, Optimization of infrastructure systems maintenance and improvement policies, Journal of Infrastructure Systems, vol. 5, American Society of Civil Engineers, Virginia, USA, 1999, pp. 124-134 34

[17] W. Yunzhong and P.S. Goldschmidt, Optimisation of Pavement Maintenance and Renewal Activities, University of Western Australia, Western Australia, 2008

[18] W.T. Chan, T.F. Fwa and K.Z. Hoque, Constraint handling methods in pavement maintenance programming, Transportation Research Part C: Emerging Technologies, Elsevier Ltd., vol. 9, issue 3, June, 2001, pp. 175-190 39

[19] T.F. Fwa, W.T. Chan, and K.Z. Hoque, Analysis of pavement management activities programming by genetic algorithms, Transportation Research Record: Journal of the Transportation Research Board, Washington DC, USA, No. 1643, 1998, p. 196

[20] G. Cancian, G. Chai and W. Pullan, Pavement Maintenance Treatment Scheduling Using a Genetic Algorithm, Accepted for presentation at the 9th International Conference on Road and Airfield Pavement Technology, Dalian China, August 2015

[21] G. Cancian, G. Chai, and W. Pullan. “Budget Constrained Pavement Maintenance Scheduling using a Parallel Genetic Algorithm”. In: The Fourth International Conference on Soft Computing Technology in Civil, Structural and Environmental Engineering. Ed. by Y. Tsompanakis, J. Kruis, and B. H. V. Topping. paper 31. Stirlingshire, United Kingdom: Civil-Comp Press, 2015. url: http://dx.doi.org/10.4203/ccp.109.31

[22] J. H. Holland, “Genetic algorithms,” Scientific American Magazine, pp. 66–72, July 1992

ISBN: 978-988-14049-1-6
ISSN: 2078-0958 (Print); ISSN: 2078-0966 (Online)