

RC6 as the AES

Ronald L. Rivest¹, M.J.B. Robshaw², and Yiqun Lisa Yin³

¹ M.I.T. Laboratory for Computer Science, 545 Technology Square,
Cambridge, MA 02139, USA. rivest@mit.edu

² 88 Hadyn Pk. Rd., London, W12 9AG, UK. mrobshaw@supanet.com

³ NTT Multimedia Communications Laboratories, 250 Cambridge Ave.,
Palo Alto, CA 94306, USA. yiqun@nttmcl.com

Introduction

After more than a year of design and nearly two years of scrutiny, the process to choose the Advanced Encryption Standard is drawing to a close. We are now left with five designs that would each be a good choice as the final AES. These five ciphers have radically different design philosophies and they have very different security and performance properties. No one cipher sticks out as being the natural choice in all respects.

During the design of RC6 our pragmatic aim was to satisfy as many goals as possible while keeping the cipher simple. Only by keeping a cipher simple can one achieve a well-understood level of security, good performance, and a versatility of design that makes the cipher highly adaptable to future demands.

We believe that we have been successful in this approach and developments over the last two years have only served to strengthen our views. We believe that RC6 would make an excellent choice as the final AES.

Security through simplicity

Despite the talk of “margins for security” and “fair” or “minimal” round assessments, the most important measure of the likely security of a cipher is quite simply the amount of scrutiny it has received. Yet it is not clear how much attention the different ciphers have received. Cryptanalysts have full-time jobs teaching in a university or working on a range of unrelated industry projects. Very few, if any, will have looked at more than two finalists in any depth, let alone all five.

A simple cipher is one that is easily described and readily remembered. It will, as a direct result, be analyzed and scrutinized widely [2, 4, 5, 8, 11]. Not only will it receive the greatest quantity of analysis - it will also receive the most accurate analysis. During the design of RC6 we performed what we believe to be one of the most accurate assessments of the security of any of the AES finalists [4]. RC6 is not so complicated that approximating models have to be introduced (as with MARS [3] and Twofish [17]). Instead we were able to get a remarkably accurate view of the strength offered by RC6 using direct analysis⁴. In this way we were

⁴ Since it is easy to define simplified and small block-size variants of RC6, the cryptanalyst can perform far more extensive analysis and experimentation.

able to make a careful decision on how many rounds RC6 should have so that we delivered good performance once our security goals had been attained. In the case of some finalists new attacks have improved on the work of the designers. Yet it is a vindication of our approach that when other techniques are applied, as was done by Knudsen and Meier [11] (and also Baudron et al. [2]), they give surprisingly similar results to those provided by our own analysis. This isn't a "small margin for security". Rather it is a carefully assessed, and remarkably accurate margin for security.

As well as being earned, some faith in a cipher can be inherited. The time for assessment of the finalists throughout the AES process has been a little less than two years. By building on the knowledge of earlier ciphers we gain insight into the security of a new cipher. Clearly RC6 was designed in the light of experience gained with perhaps the most studied modern cipher, RC5 [14]. And not only with regards to the structure of the round function. We decided to choose a key schedule for RC6 that was identical to that for RC5. No other AES finalist uses a key schedule that has been open to public analysis for nearly six years. Given the problems some finalists have in the key schedule, either with key separation in the case of Twofish [12] or with related-key attacks in the case of Rijndael [7], this is a very important attribute.

The AES effort is so important that we should not be relying on crude and subjective metrics for our decisions. The process of subtracting some arbitrary number of rounds from the number of proposed rounds - arbitrary numbers that might in one case be taken from the designers documentation and in another from direct independent analysis - can be a misleading way of comparing the AES finalists. To quote [18]: "These comparisons are fundamentally flawed, because they unfairly benefit algorithms that have been cryptanalyzed the least." Instead, the true security of a cipher depends on

- the amount of cryptanalytic scrutiny received,
- the accuracy of existing cryptanalysis,
- the ease with which verifying experiments can be conducted on a cipher,
- the amount of earlier cryptanalytic work that can be used in the assessment of the cipher, and,
- the accuracy of the designers initial estimates.

We believe that on all counts RC6 is most suited to be chosen as the AES.

Performance through simplicity

Most of today's high-end computing base is deployed in PC's either in the workplace or at home, and these are 32-bit machines. Here RC6 typically offers exemplary performance. Some restricted devices that are currently quite widely deployed are 8-bit based. These might include a relatively insignificant fraction of mobile devices, but would most likely be smart cards. However, when we couple the needs of greater processing power with the inevitable drop in prices of 32-bit processors, it is very clear that the mobile computing device market,

including smart card applications, will inevitably shift to a 32-bit oriented processor base. This trend may take a few years to come to fruition, but its results are likely to be with us for the 20 or 30 years that might be required for the AES.

With regards to very cheap smart-cards with old 8-bit processors, it has already been observed [9] that such very cheap smart-cards are vulnerable to system attacks and are inherently insecure. Such insecurities would apply to any of the AES finalists. As a result we should be careful that we do not place too much weight on the performance of a cipher in an environment that is both insecure now and obsolete (perhaps even non-existent) in a few years time. Nevertheless such processors are currently deployed and the AES may well be desired in such applications. The first question we should ask is whether performance is an important issue in such situations? What applications are going to be used on such cheap 8-bit smart cards? Certainly they won't require bulk encryption - at most a few blocks of data will be processed. So, the performance of any of the five AES finalists is going to be adequate.

On a separate issue it is repeatedly claimed (almost to the point of folklore and most surprisingly in [18]) that an implementation of RC6 requires at least 176 bytes of RAM. Yet Keating [10] has already shown that this is not the case and that RC6 can be implemented in around 120 bytes of RAM. So we can conclude that all the AES finalists can be implemented, and can be expected to offer adequate performance, on cheap (insecure) low-end smart cards.

Looking to future architectures, fine-grained estimates today of performance on future architectures really don't seem to be terribly useful. Technology evolves in unpredictable ways (for instance the growing significance of DSPs) and it seems likely that technology will evolve to best support whichever of the AES finalists is chosen. Instead, experience in the area of 32-bit processors shows that there is nothing intrinsically unsuitable about any of the five finalists for future architectures and future designs can be expected to devote significant support to providing the best possible performance from the final AES.

We provide some additional observations.

- Hand-optimized assembly code will offer the best algorithm performance on any processor. Yet often, developers will use portable code in a higher-level language and compile it for the environment of use. Under such circumstances the simplicity of a cipher is very important since it allows a compiler to produce well-optimized code. This means that good performance can be achieved without time-consuming and costly hand optimizations or lengthy code that tries to choose among a dozen different optimization strategies.
- The simplicity of a cipher is most acutely reflected in the Java performance of a cipher. This is in terms of code-size, performance, and potentially most critically, the amount of dynamic RAM used during the encryption process. With the increased importance of the Internet and its extension to mobile devices, the performance of the finalist in Java could well be vital. While there may well be many small processors in the coming years [18] many of them will in fact be Java-based, for instance in set-top boxes.

- One possible future trend is the growth of the market [13] for DSPs and/or microprocessors with DSP capability. RC6 not only performs very well on processors of this type [19], but gains its impressive performance without look-up tables which provide additional burdens on memory requirements.

We believe that excellent performance of RC6 on 32-bit processors, the close convergence in performance between simple compiled code and hand-optimized assembly, and outstanding performance in Java and in DSP environments, all make RC6 ideally suited to be chosen as the AES.

Versatility through simplicity

One of the early stated aims of the AES process was that the final cipher be “simple and versatile”. For RC6 these were design goals.

RC6 is fully parameterized; the number of encryption rounds, the size of the encryption key (not just the three must-support values of 128, 192, and 256 bits), and the block-size can all be easily and readily changed. This kind of flexibility is an integral design feature. For most of the other finalists it is not at all clear how a change to the block size, or the use of an extremely long encryption key, would be accommodated.

These could be important considerations. For some applications, a developer may wish to call on a 64-bit block cipher perhaps as a drop-in replacement to DES. With RC6 as the AES, such a variant is readily described. At the other extreme, it is possible that in the near future a 256-bit hash value will be preferred. The most natural way to do this when using an AES candidate as the basis for a hash function would be to change the block-size.

As another example of the flexibility of RC6, the key schedule allows for very long keys (for example up to 1024 bits) to be used without a compromise to performance. This is not that important for encryption, but it does provide extraordinary improvements to the performance of the Davies-Meyer hashing mode [16]; potentially to the point of providing hashing performance comparable to that offered by dedicated hash functions.

Simplicity and versatility go hand-in-hand. Once again, we believe that RC6 would be the most suited finalist to become the AES.

Conclusions

The three most important attributes of the final AES are security, performance, and versatility. With RC6 we achieve all three goals. RC6 is so simple that the full details of the cipher can be recalled at will. Through simplicity we have developed a truly versatile cipher. We have also developed a cipher that offers exceptional performance, and gives the best all-round suitability in Java with all the implications this holds for future applications. Most importantly, though, existing analysis on RC6 is not only by far the most extensive of any of the finalists, it is also the most accurate and the most detailed.

For these reasons we believe that RC6 is ideally suited to be the final AES.

References

1. R. Anderson, E. Biham, and L.R. Knudsen. Serpent: A Proposal for the Advanced Encryption Standard.
2. O. Baudron, H. Gilbert, L. Granboulan, H. Handschuh, A. Joux, P. Nguyen, F. Noilhan, D. Pointcheval, T. Pornin, G. Poupard, J. Stern and S. Vaudenay. Report on the AES candidates. In Proceedings of *The Second AES Candidate Conference*, pages 53–67. March 22-23, 1999.
3. C. Burwick, D. Coppersmith, E. D’Avignon, R. Gennaro, S. Halevi, C. Jutla, S. Matyas, L. O’Conner, M. Peyravian, D. Safford, and N. Zunic. MARS - a candidate cipher for AES. June 10, 1998.
4. S. Contini, R.L. Rivest, M.J.B. Robshaw, and Y.L. Yin. The security of RC6. Available from www.rsasecurity.com/rsalabs/aes/.
5. S. Contini, R.L. Rivest, M.J.B. Robshaw, and Y.L. Yin. Improved analysis of some simplified variants of RC6. In L. Knudsen, editor, *Fast Software Encryption, Lecture Notes in Computer Science Volume 1626*, pages 1-15, Springer-Verlag, 1999.
6. J. Daemen and V. Rijmen. AES Proposal: Rijndael. June 11, 1998.
7. N. Ferguson, J. Kelsey, B. Schneier, M. Stay, D. Wagner, and D. Whiting. Improved cryptanalysis of Rijndael. Preprint.
8. H. Gilbert, H. Handschuh, A. Joux, and S. Vaudenay. A statistical attack on RC6. Preprint.
9. S. Halevi. Suggested “tweaks” for the MARS cipher. Submitted to NIST at the end of Round 1 evaluation. Available via csrc.nist.gov.
10. G. Keating. Performance analysis of AES candidates on the 6805 CPU core. In Proceedings of *The Second AES Candidate Conference*, pages 109–114. March 22-23, 1999. Available from www.ozemail.com.au/geoffk/aes-6805.
11. L.R. Knudsen and W. Meier. Correlations in RC6. Preprint.
12. F. Mirza and S. Murphy. An observation on the key schedule of Twofish. Proceedings of the Second AES Candidate Conference, pages 151-154.
13. O. Port. Chips for the post-PC era. *Business Week*, Annual Special Issue, page 96, March 27, 2000.
14. R.L. Rivest. The RC5 encryption algorithm. In B. Preneel, editor, *Fast Software Encryption, Lecture Notes in Computer Science Volume 1008*, pages 86-96, Springer-Verlag, 1995. Available from theory.lcs.mit.edu/~rivest/.
15. R.L. Rivest, M.J.B. Robshaw, R. Sidney, and Y.L. Yin. The RC6 Block Cipher. v1.1, August 20, 1998. Available from www.rsasecurity.com/rsalabs/aes/
16. M.J.B. Robshaw. Hashing with the AES finalists. Preprint.
17. B. Schneier, J. Kelsey, D. Whiting, D. Wagner, C. Hall, and N. Ferguson. Twofish: A 128-bit Block Cipher. 15 June, 1998.
18. B. Schneier and D. Whiting. A performance comparison of the five AES finalists. Preprint.
19. T. Wollinger, M. Wang, J. Guajardo, and C. Paar. How well are high-end DSPs suited for the AES algorithms? Preprint.