TinyECC: A Configurable Library for Elliptic Curve Cryptography in Wireless Sensor Networks*

An Liu
Department of Computer Science
NC State University, Raleigh, NC 27695
email: aliu3@ncsu.edu

Peng Ning
Department of Computer Science
NC State University, Raleigh, NC 27695
email: pning@ncsu.edu

Abstract

Public Key Cryptography (PKC) has been the enabling technology underlying many security services and protocols in traditional networks such as the Internet. In the context of wireless sensor networks, elliptic curve cryptography (ECC), one of the most efficient types of PKC, is being investigated to provide PKC support in sensor network applications so that the existing PKC-based solutions can be exploited.

This paper presents the design, implementation, and evaluation of TinyECC, a configurable *library for ECC operations in wireless sensor networks. The primary objective of TinyECC is to provide a* ready-to-use, publicly available *software package for ECC-based PKC operations that can be* flexibly configured and integrated *into sensor network applications. TinyECC provides a number of optimization switches, which can turn specific optimizations on or off based on developers' needs. Different combinations of the optimizations have different execution time and resource consumptions, giving developers great flexibility in integrating TinyECC into sensor network applications. This paper also reports the experimental evaluation of TinyECC on several common sensor platforms, including MICAz, Tmote Sky, and Imote2. The evaluation results show the impacts of individual optimizations on the execution time and resource consumptions, and give the most computationally efficient and the most storage efficient configuration of TinyECC.*

1. Introduction

Recent technological advances have made it possible to develop wireless sensor networks consisting of a large number of low-cost, low-power, and multi-functional sensor nodes that communicate over short distances through wireless links. Such sensor networks are ideal candidates for a wide range of applications such as monitoring of critical infrastructures, data acquisition in hazardous environments, and military operations. The desirable features of wireless sensor networks have attracted many researchers to develop protocols and algorithms that can fulfill the requirements of these applications.

Security services such as authentication and key management are critical to communication security in wireless sensor networks as well as the security of sensor network applications. In traditional networks such as the Internet, Public Key Cryptography (PKC) has been the enabling technology underlying many security services and protocols (e.g., SSL [3] and IPsec [18]). For example, PKC has been used to bootstrap symmetric session keys and authenticate messages to multiple receivers. However, in wireless sensor networks, PKC has not been widely adopted due to the resource constraints on sensor platforms, particularly the limited and depleteable battery power.

There has been intensive research aimed at developing techniques that can bypass PKC operations in sensor network applications. For example, there has been a substantial amount of research on random key pre-distribution for pairwise key establishment (e.g., [13, 23]) and broadcast authentication (e.g., [24, 25]). However, these alternative approaches do not offer the same degree of security or functionality as PKC. For instance, none of the random key pre-distribution schemes can guarantee key establishment between any two nodes and tolerate arbitrary node compromises at the same time. As another example, the aforementioned broadcast authentication schemes, which are all based on TESLA [32], require loose time synchronization, which itself is a challenging task to achieve in wireless sensor networks. In contrast, PKC can address all these problems easily. Pairwise key establishment can always be achieved using, for example, the Diffie-Hellman (DH) key exchange protocol [12], without suffering from the node compromise problem. Similarly, broadcast authentication can be provided with, for example, the ECDSA digital sig-

*This work is supported by the National Science Foundation under grants CAREER-0447761 and CNS-0721424, and by the Army Research Office under grants W911NF-05-1-0247 and W911NF-04-D-0003. The contents of this paper do not necessarily reflect the position or the policies of the U.S. Government.

nature scheme [7], without requiring time synchronization. Thus, it is desirable to explore the application of PKC on resource constrained sensor platforms.

There have been a few recent attempts to use PKC in wireless sensor networks [15, 26, 33], which demonstrate that it is feasible to perform limited PKC operations on the current sensor platforms such as MICAz motes [2]. Elliptic Curve Cryptography (ECC) has been the top choice among various PKC options due to its fast computation, small key size, and compact signatures. For example, to provide equivalent security to 1024-bit RSA, an ECC scheme only needs 160 bits on various parameters, such as 160-bit finite field operations and 160-bit key size [8].

Despite the recent progress on ECC implementations on sensor platforms, all the previous attempts [15, 26, 33] have limitations. In particular, all these attempts were developed as independent packages/applications without seriously considering the resource demands of sensor network applications. As a result, developers may find it difficult, and sometimes impossible, to integrate an ECC implementation with the sensor network applications, though the ECC implementation may be okay on its own. For example, an ECC implementation may require so much RAM that it would be impossible to fit both the sensor network application and the ECC implementation on the same node.

Moreover, various optimization techniques are available to speed up the ECC operations. Such optimizations, however, typically will increase the ROM and RAM consumptions, though they may reduce the execution time and energy consumption. It is not clear what optimizations should be used and how they should be combined to achieve the best trade-off among security protection, computation overheads, and storage requirements. Additional research is necessary to clarify these issues and facilitate the adoption of ECC-based PKC in wireless sensor networks.

It is certainly possible to have dedicated PKC hardware included on sensor platforms. However, given that there is no PKC hardware that is currently available on current sensor platforms, it is a sensible choice to explore software approaches for PKC support on sensor platforms.

In this paper, we present the design, implementation, and evaluation of TinyECC, a *configurable* library for ECC operations in wireless sensor networks.[1] The primary objective of TinyECC is to provide a *ready-to-use, publicly available* software package for ECC-based PKC operations that can be *flexibly configured and integrated* into sensor network applications.

Targeted at TinyOS [5], TinyECC is written in nesC [14], with occasional in-line assembly code to achieve further speedup for popular sensor platforms including MICAz [2], TelosB [4], Tmote Sky [6], and Imote2 [1]. A unique feature of TinyECC is its *configurability*. TinyECC includes almost all known optimizations for ECC operations. Each optimization is controlled by a software switch, which can turn the optimization on or off based on developers' needs. Different combinations of optimizations have different ROM/RAM consumption, execution time, and energy consumption. This gives the developers great flexibility in integrating TinyECC in their applications.

To provide guidance in using TinyECC, we perform a series of experiments with different combinations of activated optimizations. To understand the impact of each optimization technique, we compare the execution time, ROM/RAM consumption, and energy consumption with and without the given optimization enabled on MICAz [2], Tmote Sky [6], and Imote2 [1]. In addition, our experiments also present the performance results and the resource usages for the most computationally efficient configuration (i.e., fastest execution and least energy consumption) and the most storage-efficient configuration (i.e., least ROM and RAM usage) of TinyECC on these common sensor platforms, respectively.

The contribution of this paper is two-fold: First, we develop TinyECC to allow flexible integration of ECC-based PKC in sensor network applications. Second, we perform a substantial amount of experimental evaluation using representative sensor platforms, including MICAz [2], TelosB [4], Tmote Sky [6], and Imote2 [1]. The experimental results provide useful experience and guidance for developers to choose different TinyECC optimizations for their needs.

The remainder of this paper is organized as follows. Section 2 discusses the design principles of TinyECC. Section 3 gives background information on ECC. Section 4 describes the optimization techniques adopted by TinyECC. Section 5 discusses the implementation of TinyECC. Section 6 presents the experimental evaluation of TinyECC on MICAz, Tmote Sky, and Imote2. Section 7 discusses the related work, and Section 8 concludes this paper.

2. Design Principles

As mentioned earlier, the primary objective of TinyECC is to provide a *ready-to-use, publicly available* software package for ECC-based PKC operations that can be *flexibly configured and integrated* into sensor network applications. To make sure we achieve this objective, we follow several principles in the design and development of TinyECC.

Security: TinyECC should provide PKC schemes that have proven to be secure. To follow this principle, TinyECC only includes support for the well-studied ECC schemes such as ECDSA, ECDH, and ECIES, which are defined in the Standards for Efficient Cryptography [8]. Moreover, TinyECC also includes elliptic curve parameters recommended by SECG (Stands for Efficient Cryp-

[1] TinyECC 1.0 and its previous versions are publicly available at `http://discovery.csc.ncsu.edu/software/TinyECC/`.

tography Group), such as `secp160k1`, `secp160r1` and `secp160r2`, as defined in [9].

Portability: TinyECC should run on as many sensor platforms as possible. Due to this reason, we choose to implement TinyECC on TinyOS [5], which is a popular, open-source OS for networked sensors. All the TinyECC components have nesC [14] implementations, though some modules also include inline assembly code, which can be turned on for faster execution on some sensor platforms. This allows TinyECC to be compiled and used on any sensor platform that can run TinyOS. TinyECC has been tested successfully on MICAz, TelosB, Tmote Sky, and Imote2.

Resource Awareness and Configurability: TinyECC should accommodate the typical resource constraints on sensor nodes. Moreover, TinyECC should allow for flexible configuration so that it can take advantage of the available resources on a wide spectrum of sensor platforms. To follow this principle, TinyECC is implemented carefully to avoid unnecessary resource usage. Moreover, TinyECC uses a set of optimization switches, which can be turned on or off to achieve different combinations of performance and resource consumptions.

Efficiency: TinyECC should be computationally efficient to reduce the battery consumption as well as the delay introduced by PKC operations. We make three design decisions to improve the efficiency of TinyECC. The first is about the type of finite fields over which the ECC operations are performed. ECC can be implemented over either a prime field F_p, where p is a large prime number, or a binary extension field F_{2^m}, where m is an integer. Since arithmetic operations over F_{2^m} are insufficiently supported by micro-controllers, we choose to support prime fields F_p in TinyECC. Second, we adopt almost all existing optimizations for ECC operations in TinyECC. As mentioned earlier, these optimizations can be turned on or off to balance the efficiency and the resource requirements. Third, we include inline assembly code in critical parts of TinyECC for popular sensor platforms, including MICAz, TelosB, Tmote Sky, and Imote2.

Functionality: TinyECC should support the typical demands for PKC. To follow this principle, the current version of TinyECC includes a digital signature scheme (ECDSA), a key exchange protocol (ECDH), and a public key encryption scheme (ECIES). These cover all typical uses of PKC.

3. Background on ECC

In this and next sections, we give an overview of ECC and the optimizations adopted by TinyECC as a convenient reference. The reader can find details in the references.

Elliptic curve cryptography (ECC) is an approach to public-key cryptography based on the algebraic structure of elliptic curves over finite fields [16]. Elliptic curves used in cryptography are typically defined over two types of finite fields: prime fields F_p, where p is a large prime number, and binary extension fields F_{2^m}. For space reasons, we focus on elliptic curves over F_p in this paper.

An elliptic curve over F_p is defined by a cubic equation $y^2 = x^3 + ax + b$, where $a, b \in F_p$ are constants such that $4a^3 + 27b^3 \neq 0$ [16]. An elliptic curve over F_p consists of the set of all pairs of affine coordinates (x, y) for $x, y \in F_q$ that satisfy an equation of the above form and an infinity point $\mathcal{O}$. The points on an elliptic curve form an abelian group with $\mathcal{O}$ as the additive identity. (The formulas defining point addition and its special case, point doubling, can be found in [16].)

For any point G on an elliptic curve, the set $\{\mathcal{O}, G, 2G, 3G, ...\}$ is a cyclic group [16]. The calculation of kG, where k is an integer, is called a *scalar multiplication*. The problem of finding k given points kG and G is called the *elliptic curve discrete logarithm problem (ECDLP)*. It is computationally infeasible to solve ECDLP for appropriate parameters [16]. The hardness of ECDLP allows several cryptographic schemes based on elliptic curves.

TinyECC includes three well-known ECC schemes: (1) the Elliptic Curve Diffie-Hellman (ECDH) key agreement scheme, (2) the Elliptic Curve Digital Signature Algorithm (ECDSA), and (3) the Elliptic Curve Integrated Encryption Scheme (ECIES). ECDH is a variant of the Diffie-Hellman key agreement protocol [12] on elliptic curve groups. ECDSA is a variant of the Digital Signature Algorithm (DSA) [29] that operates on elliptic curve groups. ECIES is a public-key encryption scheme which provides semantic security against an adversary who is allowed to use chosen-plaintext and chosen-ciphertext attacks [16]. ECIES is also known as the Elliptic Curve Augmented Encryption Scheme (ECAES) or simply the Elliptic Curve Encryption Scheme. These ECC schemes allow smaller key sizes for similar security level to the alternatives such as the original DH and DSA schemes. For each of the schemes, a party that would like to use the scheme needs to agree on some domain parameters such as the elliptic curve and a point G on the curve, and must have a key pair consisting of a private key d and a public key $Q = dG$. The specification of ECDH, ECDSA, and ECIES can be found in [8, 16].

4. Optimizations Adopted by TinyECC

In this section, we briefly discuss the optimization techniques adopted by TinyECC. We will omit the details, since the focus of this paper is not these individual optimization techniques. More information about these techniques can be found in the relevant references.

4.1. Optimizations for Large Integer Operations

Barrett Reduction [28]: A straightforward way to perform large integer modular reductions is to use division [19]. A nice side effect is that it reuses the code of division, thus resulting in more compact code size.

Barrett Reduction is an alternative method for modular reduction [28]. It converts the reduction modulo an arbitrary integer to two multiplications and a few reductions modulo integers of the form 2^n. When used to reduce a single number, Barrett reduction is slower than a normal division algorithm. However, when used to reduce various numbers modulo the same number many times, by precomputing some value, Barrett reduction can achieve faster speed than modular reductions obtained by division. Details of Barrett reduction can be found in [28].

In TinyECC, since almost all the modular operations are modulo the same prime number p, Barrett reduction can potentially speed up the computation. However, this requires the implementation of a separate reduction algorithm, which implies larger code size (i.e., greater ROM requirement) on sensor nodes. In addition, Barrett reduction also increases RAM use. Assume the target microcontroller has a w-bit word size. Given a finite field F_p, where p is a k words long prime number, Barrett reduction requires the pre-computation of $\mu = \lfloor \frac{b^k}{p} \rfloor$, where $b = 2^w$ (e.g., $b = 2^8$ on a 8-bit processor). This number μ has to be stored and used throughout all the modular reductions. Thus, to exchange for faster computation, Barrett reduction requires more ROM and RAM than the traditional division-based modular reduction.

Hybrid Multiplication and Hybrid Squaring [15]: Standard large integer multiplication algorithms [19] store the operands and the product in arrays. When such an algorithm is implemented in a high-level language such as nesC, the compiler cannot use the registers in the microcontroller efficiently, and the binary code usually needs to load the operands from memory to registers multiple times [15]. Gura et al. [15] proposed a hybrid multiplication algorithm, which was intended for assembly code. This algorithm can maximize the utilization of registers and reduce the number of memory operations. TinyECC adopts this hybrid multiplication algorithm for MICAz [2], TelosB [4]/Tmote Sky [6], and Imote2 [1]. Indeed, the code can be used on any sensor platforms that have processors using the same instruction sets.

In addition to hybrid multiplication, we also customize the hybrid multiplication algorithm for squaring operations by using the fact that the two multiplicative operands in squaring are the same. This further reduces the execution time for squaring at the cost of larger code size.

4.2. Optimizations for ECC Operations

Projective Coordinate Systems [16]: As discussed earlier, an elliptic curve consists of the infinity point $\mathcal{O}$ and the set of points in the affine coordinates (x,y) for $x,y \in F_p$ that satisfy the defining equation. Alternatively, a point on an elliptic curve can be represented in a projective coordinate system in the form of (x,y,z).

Point addition and point doubling are critical operations in ECC, which are building blocks for scalar multiplications required by all ECC schemes. These operations in affine coordinate system require modular inversion operations, which are much more expensive than other operations such as modular multiplications. Using a projective coordinate system [16], modular inversions can be removed with the compensation of a few modular multiplications and squares. As a result, the execution times of point addition and point doubling based on projective coordinate system are faster than those based on affine coordinate system, respectively [16].

TinyECC uses two additional optimizations along with projective coordinate representation, which can further reduce both the execution time and the program size. The first is a *mixed point addition algorithm* [16], which adds a point in projective coordinate and a second point in affine coordinate. This algorithm can be used in scalar multiplications to further reduce the number of modular multiplications and squares, leading to smaller and faster code. The second is *repeated Doubling* [16] for scalar multiplication. If consecutive point doublings are to be performed, the repeated doubling algorithm may be used to achieve faster performance than repeated use of the doubling formula. In m consecutive doublings, this algorithm trades $m-1$ field additions, $m-1$ divisions by two, and a multiplication for two field squarings (in comparison with repeated applications of the plain point doubling algorithm) [16].

Though reducing the execution time, the projective coordinate representation requires a larger code size (for more complex formula) and more RAM (for storing additional variables) than the affine coordinate system.

Sliding Window for Scalar Multiplications [16]: Scalar multiplication is a basic operation used by all ECC schemes. It is in the form of kP, where k is an integer and P is a point on an elliptic curve. In the most straightforward method to compute kP, k is scanned from the most significant bit to the least significant bit. When each bit is scanned, the algorithm needs to compute a point doubling. When the scanned bit is "1", the algorithm also needs to perform a point addition. The sliding window method can speed up the scalar multiplication by scanning w bits at a time. Each time when a w-bit window is scanned, the algorithm needs to perform w point doublings. By precomputing $2P$, $3P$, ..., and $(2^w-1)P$, the sliding window method

only needs to perform 1 point addition every w bits, and thus has less computational cost.

It is easy to see that the sliding window method will increase both the ROM (for additional code size) and RAM (for storing the pre-computed points) consumptions.

Shamir's Trick [16]: This optimization is only used for the verification of ECDSA signatures. The verification of ECDSA signature requires the computation of the form $aP + bQ$, where a, b are integers and P, Q are two points on an elliptic curve. A straightforward implementation requires two scalar multiplications and a point addition. However, Shamir's trick allows us to compute the above value at a cost close to one scalar multiplication. Specifically, with pre-computed $P + Q$, we may scan the (same) bits of a and b from the most significant one to the least significant one. For each bit, we need double the intermediate value, which is initialized as the infinity point. If the scanned bit positions are $\langle a_i = 0, b_i = 1\rangle$, $\langle a_i = 1, b_i = 0\rangle$, or $\langle a_i = 1, b_i = 1\rangle$, we add P, Q, or $P + Q$ to the intermediate value. This reduces two scalar multiplications to be a bit more expensive than one such operation.

Similar to the sliding window method, Shamir's trick will increase both the ROM (for additional code size) and RAM (for storing the pre-computed $P + Q$) consumptions.

Curve Specific Optimization [15]: A number of elliptic curves specified by NIST [30] and SECG [9] use pseudo-Mersenne primes. A pseudo-Mersenne prime is of the form $p = 2^n - c$, where $c \ll 2^n$. Reduction modulo a pseudo-Mersenne prime can be performed by a few modular multiplications and additions without any division operation. As a result, the time for modular reduction can be reduced significantly. Thus, using elliptic curves over a pseudo-Mersenne prime can achieve additional performance gain.

5. Implementation

We implemented TinyECC on TinyOS [5], an open source operating system designed for wireless embedded sensor networks. The current version of TinyECC provides support for ECDSA (digital signatures), ECDH (pairwise key establishment), and ECIES (PKC-based encryption). Most of the code was written in nesC [14] for portability reasons. To best harness the capabilities of the processors on popular sensor platforms such as MICAz and TelosB, we also provided inline assembly implementation of some critical operations, such as large integer multiplications.

To save implementation efforts, we ported the C code of large integer operations in RSAREF 2.0 [20] to nesC code on TinyOS. These include modular addition, subtraction, multiplication, division, inverse, and exponentiation operations. We then implemented all the elliptic curve operations and the optimization techniques discussed earlier.

TinyECC has been released publicly at `http://discovery.csc.ncsu.edu/software/TinyECC/`. Some preliminary versions have been adopted by other researchers (e.g., [11, 21, 27]). As discussed earlier, starting from the current version, we added a set of optimization switches to provide for flexible configuration of TinyECC so that it can be integrated into sensor network applications with different resource consumptions and performance demands.

Table 1 lists the optimization switches available in the current version of TinyECC. All optimization switches can be turned on or off by a simple configuration at compile time, or slight modification in the source code. Moreover, when the sliding window method is used, an additional parameter defining the size of the window (e.g., $w = 4$) needs to be specified.

6. Evaluation

We performed a series of experiments to evaluate TinyECC on four representative sensor platforms, including MICAz [2], TelosB [4], Tmote Sky [6], and Imote2 [1].

The objective of these experiments is three-fold: First, we want to measure the performance and resource consumption of TinyECC on a spectrum of sensor platforms, ranging from the low-end ones (such as MICAz, TelosB, and Tmote Sky) to high-end ones (such as Imote2). Second, we would like to understand the impact of the optimizations adopted by TinyECC on performance and resource consumption. Finally, we would like to provide detailed performance results and resource demands for commonly desirable configurations, including the configuration that provides the fastest execution time and the configuration that requires the least memory consumption. The former has the least energy consumption, while the latter is the easiest one to integrate into sensor network applications.

6.1. Methodology and Experimental Setup

Evaluation Methodology: Given seven optimization switches, four sensor platforms, where Imote2 has multiple CPU frequencies due to dynamic voltage scaling, many possible elliptic curves, and three ECC-based PKC schemes, there are a large number of experiments to perform if we have to observe the performance and resource consumptions in all cases.

To simplify the scenarios, we adopted the following methodology in our experiments. For each optimization switch, we performed two sets of experiments, referred to as *case A* and *case B*, respectively. In case A, for each optimization, we disabled all the other optimizations, and then obtained the performance and resource consumption metrics when the given optimization was enabled and disabled, respectively. In case B, we enabled all the other optimizations and obtained the evaluation metrics again when the given optimization was enabled and disabled, respectively.

Table 1. TinyECC Optimization Switches

Method	Optimization Switch	Category	Description
Barrett Reduction	BARRETT	large number	Allow Barrett reduction.
Hybrid Multiplication	HYBRID_MULT	large number	Allow hybrid multiplication in inline assembly.
Hybrid Squaring	HYBRID_SQR	large number	Allow hybrid squaring in inline assembly.
Projective Coordinate System	PROJECTIVE	EC	Use projective coordinate system along with mixed point addition and repeated doubling.
Sliding Window Method	SLIDING_WIN	EC	Use sliding window method for scalar multiplication. A window size (e.g. $w = 4$) has to be defined along with this switch.
Shamir's Trick	SHAMIR_TRICK	EC	Allow Shamir's trick when verifying ECDSA signatures. A window size (e.g. $w = 2$) has to be defined along with this switch.
Curve-Specific Optimization	CURVE_OPT	EC	Allow curve specification optimization. This has to be used for the curves defined over pseudo-Mersenne primes [9, 30].

The differences in these metrics reflect the impact of the given optimization technique.

Moreover, as discussed earlier, we also performed additional experiments to examine in detail two commonly desirable configurations: the one that provides the fastest execution time, and the one that requires the least storage.

Experimental Setup: We evaluated TinyECC on the latest CVS version of TinyOS 1.x [5]. As discussed earlier, we chose four representative sensor platforms, MICAz, TelosB, Tmote Sky, and Imote2, for the experiments, since they are popular sensor platforms and cover the 8-bit, 16-bit and 32-bit processors. Other sensor platforms (e.g., Mica2, Mica2Dot) are expected to perform similarly to one of these platforms, due to the use of the same processor.

TelosB and Tmote Sky have almost the same hardware. The only difference is that TelosB can only run at 4 MHz, while Tmote Sky can run at 8 MHz when an external resistor is enabled. We configure Tmote Sky to run at 8 MHz in our experiments. Due to the similarity between TelosB and Tmote Sky, we only report the results on Tmote Sky in this paper. The reader may refer to the technical report version of this paper [22] for experimental results on TelosB.

As a high-end sensor platform, Imote2 uses an XScale processor and supports dynamic voltage scaling. To obtain a relatively complete view of Imote2, we used four different frequencies on Imote2 in our experiments: 13MHz, 104MHz, 208MHz, and 416MHz.

By default, TinyECC includes all 128-bit, 160-bit and 192-bit ECC parameters recommended by SECG [9]. It is well-known that 160-bit ECC has the same security level as 1024-bit RSA. We selected a 160-bit elliptic curve `secp160r1` [9] to evaluate the impact of individual optimization techniques. Note that the actual selection of curves depends on the security needs in the sensor network applications, and is outside of the scope of this paper.

We used the following evaluation metrics in all experiments: ROM consumption (byte), RAM consumption (byte), execution time (ms), and energy consumption (millijoule). We used the `check_size.pl` script in the TinyOS distribution to obtain the ROM and RAM sizes required by the TinyECC components. The execution time was measured directly on the sensor nodes. To get the overall performance result, we randomly generated the parameters (e.g., random message, random public and private key pairs) other than those defining the curves, and obtained the execution time for each data point by taking the average of 10 test instances. The energy consumption was then calculated as $U \times I \times t$ based on the execution time (t), the voltage (U), and current draw (I) on these sensor platforms [1, 2, 4, 6].

6.2. Evaluation Results

Due to the space limit, we can only report a portion of the evaluation results. Please refer to the full version of this paper [22] for more details.

6.2.1. Impact of Individual Optimizations

In the following, we use the experimental results for ECDSA to show the impact of individual optimization techniques. More results on the impacts of these optimizations on ECDH and ECIES can be found in [22].

There are three aspects of the execution time for ECDSA. Figures 1(a) and 1(b) show the initialization time required to prepare for ECDSA in cases A and B, respectively. Figures 1(c) and 1(d) show the signature generation time in cases A and B, respectively. Figures 1(e) and 1(f) show the signature verification time in cases A and B, respectively.

In the initialization of ECDSA, TinyECC needs to precompute μ for Barrett reduction, a few points for the sliding window method, and a few points for Shamir's trick. In case A, as Figure 1(a) shows, only these 3 optimization techniques have impact on the initialization time. For MICAz, the initialization of the sliding window method with window size 4 requires 3,587 ms, which is longer than Shamir's trick (1,672 ms for window size 2) and Barrett reduction (6 ms). The same situation applies to TelosB/Tmote Sky, and Imote2. If we disable all these three techniques, the initialization time of ECDSA is close to 0. In case B, as Figure 1(b) shows, the disabling of selected optimization technique does not reduce the initialization time dramatically, except that the disabling of the sliding window method reduces the initialization time by half.

In Figure 1, we can see that *PROJECTIVE* is the most

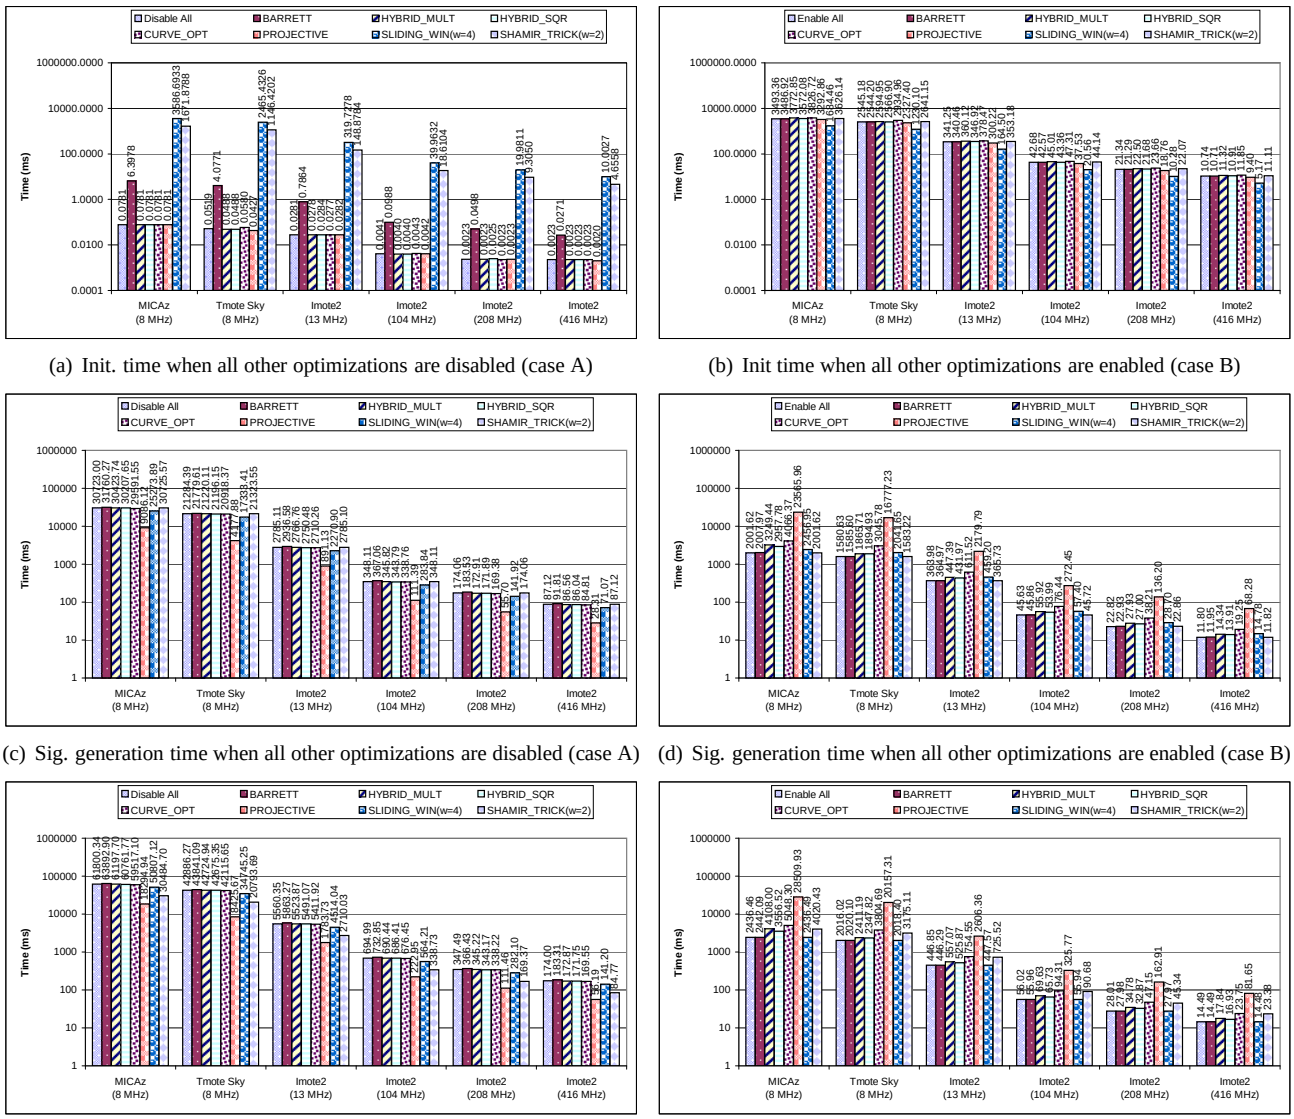

(a) Init. time when all other optimizations are disabled (case A)

(b) Init time when all other optimizations are enabled (case B)

(c) Sig. generation time when all other optimizations are disabled (case A)

(d) Sig. generation time when all other optimizations are enabled (case B)

(e) Sig. verification time when all other optimizations are disabled (case A)

(f) Sig. verification time when all other optimizations are enabled (case B)

Figure 1. ECDSA timing result (Projective coordinate system is the most effective optimization, while Barrett reduction does not have obvious impact.)

effective switch to improve the speed of signature generation and verification. In case A, by enabling the *PROJECTIVE* switch, the signature generation and verification of all platforms can speed up by at least 3 times. In case B, if we disable the *PROJECTIVE* switch, the signature generation and verification has at least 6 times slowdown compared with enabling all optimization techniques.

Although *PROJECTIVE* is the most efficient switch, it increases the ROM usage. Figures 2(a) and 2(b) show that when the *PROJECTIVE* switch is enabled in case A, the ROM size is increased by 1,218, 1,326, and 1,752 bytes for MICAz, TelosB/Tmote Sky, and Imote2, respectively, while the RAM size does not change at all. In case B, as Figures 2(c) and 2(d) show, disabling the *PROJECTIVE* switch can save 3,816, 3,880, and 4,660 bytes in ROM for MICAz, TelosB/Tmote Sky, and Imote2, respectively. The *PROJECTIVE* switch is the most effective switch to speed up ECDSA operations, but it also incurs larger ROM consumption than any other optimization technique.

SHAMIR_TRICK is also an efficient option to speed up ECDSA signature verification. From Figure 1(e), we can see that the verification can speed up by 2 times on all platforms when enabling *SHAMIR_TRICK* in case A. Both ROM and RAM sizes are increased. In case A, the RAM size is increased by 634, 676, and 784 bytes for MICAz, TelosB/Tmote Sky, and Imote2, respectively. Similarly, the

ROM size of MICAz, TelosB/Tmote Sky and Imote2 is increased by 638, 632, and 620 bytes, respectively. In case B, disabling *SHAMIR_TRICK* makes verification 1.6 times slower but saves 2,148, 2,068, and 2,208 bytes in ROM for MICAz, TelosB/Tmote Sky, and Imote2, respectively. The RAM size does not decrease much because the sliding window method is used for verification when *SHAMIR_TRICK* is disabled.

Now let us take a look at the *SLIDING_WIN* option. In case A, as Figures 1(c), 1(e), 2(a) and 2(b) show, enabling *SLIDING_WIN* can make signature generation and verification 1.2 times faster at the cost of dramatic RAM increase (1,262, 1,328 and 1,472 bytes for MICAz, TelosB/Tmote Sky, and Imote2, respectively). In case B, as Figures 1(d), 1(f), 2(c) and 2(d) show, disabling *SLIDING_WIN* can save 632, 668, and 752 bytes of RAM for MICAz, TelosB/Tmote Sky, and Imote2 with 1.2 times slower signature generation and verification. Since MICAz and TelosB/Tmote Sky are low-end sensor platforms, they have much smaller RAM (4kB, 10kB) compared with Imote2 (256kB). Before enabling *SLIDING_WIN*, we should be very careful if the sensing application has large RAM consumption. Since *SLIDING_WIN* is the most RAM consuming switch in TinyECC, application developers may disable it or reduce the window size to reserve more RAM for the applications.

Now consider the *HYBRID_MULT*, *HYBRID_SQR*, and *CURVE_OPT* options. In case A, *HYBRID_MULT*, *HYBRID_SQR* and *CURVE_OPT* do not have big impact on the timing result. However, in case B, *HYBRID_MULT* can speed up signature generation by 1.6 times for MICAz, 1.2 times for TelosB/Tmote Sky, and 1.2 times for Imote2. Similarly, it can speed up signature verification by 1.7 times for MICAz, 1.2 times for TelosB/Tmote Sky, and 1.2 times for Imote2. *HYBRID_SQR* can speed up signature generation by 1.5 times for MICAz, 1.2 times for TelosB/Tmote Sky, and 1.2 times for Imote2, and speed up signature verification by 1.5 times for MICAz, 1.2 times for TelosB/Tmote Sky, and 1.2 times for Imote2. *CURVE_OPT* can speed up signature generation by 2 times for MICAz, 1.9 times for TelosB/Tmote Sky, and 1.7 times for Imote2. Similarly, it can speed up signature verification by 2.1 times for MICAz, 1.9 times for TelosB/Tmote Sky, and 1.7 times for Imote2. The reason that *HYBRID_MULT*, *HYBRID_SQR* and *CURVE_OPT* cannot speed up ECDSA much in case A is that the *PROJECTIVE* option is disabled when each of these switches is enabled. Thus, inverse operation is the major computation of signature generation and verification. In case B, when *PROJECTIVE* is enabled, multiplication and squaring become the major computation in ECDSA.

Based on the timing results obtained for ECDSA, the effectiveness of these optimization switches in terms of

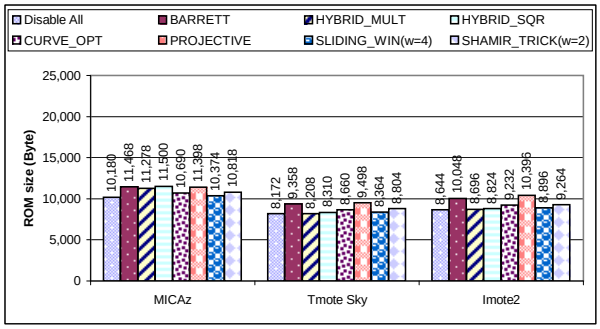

(a) ROM size w/ all other optimizations disabled (case A)

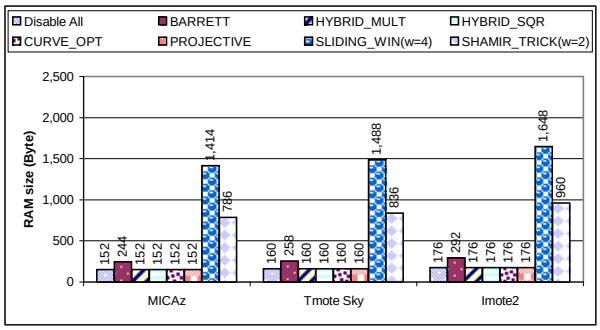

(b) RAM size w/ all other optimizations disabled (case A)

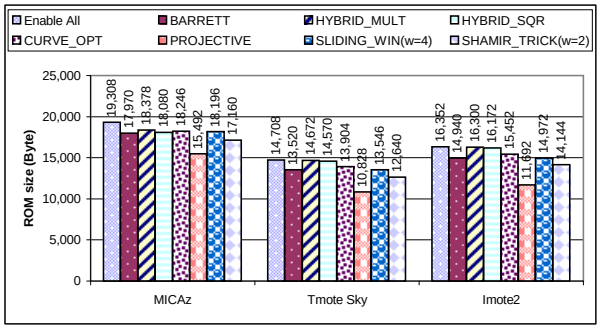

(c) ROM size w/ all other optimizations enabled (case B)

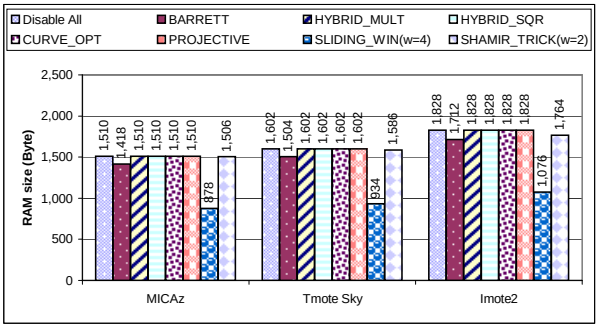

(d) RAM size w/ all other optimizations enabled (case B)

Figure 2. Code size of ECDSA (Sliding window has the most memory demand, Shamir's trick ranks the next, while all the other optimizations have similar memory demands.)

Available ROM (byte)	How to Conf. TinyECC's Switches
$[10,180,10,374)$	disable all
$[10,374,11,398)$	enable SLIDING_WIN
$[11,398,17,160)$	enable PROJECTIVE
$[17,160,17,970)$	enable all & disable SHAMIR_TRICK
$[17,970,19,308)$	enable all & disable BARRETT
$[19,308,+\infty)$	enable all

Table 2. Configuration based on Free ROM for MICAz

Available RAM (byte)	How to Conf. TinyECC's Switches
$[152,786)$	disable all
$[786,878)$	enable SHAMIR_TRICK
$[878,1,418)$	enable all & disable SLIDING_WIN
$[1,418,1,510)$	enable all & disable BARRETT
$[1,510,+\infty)$	enable all

Table 3. Configuration based on Free RAM for MICAz

execution time can be ordered as follows: *PROJECTIVE* $>$ *CURVE_OPT* $>$ *HYBRID_MULT* $>$ *HYBRID_SQR* $>$ *SLIDING_WIN* $>$ *SHAMIR_TRICK* $>$ *BARRETT*. In terms of RAM size, the optimization switches can be ordered as follows: *SLIDING_WIN* $>$ *SHAMIR_TRICK* $>$ *BARRETT* $>$ *HYBRID_MULT* $=$ *HYBRID_SQR* $=$ *CURVE_OPT* $=$ *PROJECTIVE*.

In terms of ROM size, the optimization switches are ordered differently for different platforms. For MICAz, *PROJECTIVE* $>$ *BARRETT* $\approx$ *HYBRID_SQR* $>$ *CURVE_OPT* $\approx$ *SHAMIR_TRICK* $\approx$ *HYBRID_MULT* $>$ *SLIDING_WIN*. For TelosB/Tmote Sky, *PROJECTIVE* $>$ *BARRETT* $\approx$ *SHAMIR_TRICK* $>$ *CURVE_OPT* $\approx$ *SLIDING_WIN* $>$ *HYBRID_SQR* $>$ *HYBRID_MULT*. For Imote2, *PROJECTIVE* $>$ *BARRETT* $>$ *SHAMIR_TRICK* $\geq$ *CURVE_OPT* $>$ *SLIDING_WIN* $>$ *HYBRID_SQR* $>$ *HYBRID_MULT*.

Configuration Guideline: To summarize and assist users of TinyECC, in Tables 2 and 3, we show how to choose optimization switches on MICAz motes when the amount of available ROM and RAM for TinyECC can be estimated. TinyECC requires at least 10,180 bytes ROM and 152 bytes RAM to be used on MICAz. As more ROM and RAM are available, we can gradually enable the optimization switches as indicated in these tables to get better performance. For example, having a ROM size between 11,398 bytes and 17,160 bytes allows us to enable both SLIDING_WIN and PROJECTIVE switches (but not others). Note that an optimization can be enabled if both ROM and RAM sizes allow it. Optimization switches can be determined for other platforms similarly.

6.2.2. Most Computationally Efficient Configuration

Now consider the most computationally efficient configuration. Apparently, TinyECC provides the most computationally efficient configuration when all the optimization switches are enabled. Figure 3 shows the execution time required by ECDSA initialization, signature generation, signature verification; ECIES initialization, encryption, decryption; ECDH initialization, key establishment.

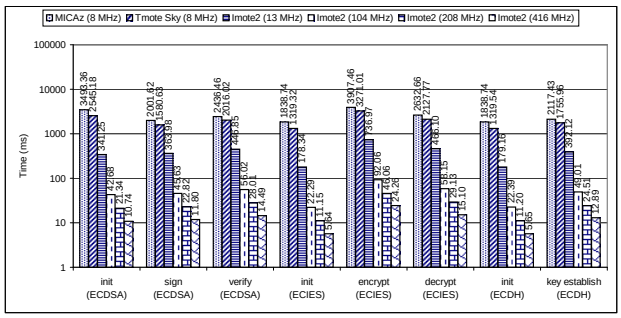

Figure 3. Execution time of ECDSA, ECIES, and ECDH w/ all optimization switches enabled

From figure 3, we can see that enabling all optimization switches requires long pre-computation. For example, it takes MICAz 3,493, 1,839 and 1,839 ms to do pre-computation for ECDSA, ECIES and ECDH, respectively. Most of the pre-computation time is for the sliding window method and Shamir's trick (ECDSA only). Tmote Sky runs slightly faster than MICAz. Running at 13 MHz, the default CPU frequency for Imote2, Imote2 is faster than Tmote Sky in all operations. If we set the frequency to 416 MHz, it only takes 12 and 14 ms to generate ECDSA signature and verify it. Moreover, it can perform ECIES encryption in 24 ms and decrypt in 15 ms. Finally, ECDH key establishment only takes 13 ms.

Enabling all optimization switches requires the largest ROM and RAM consumptions. Figure 4 shows the ROM and RAM requirements by all schemes. Imote2 has the largest RAM size due to its word size. MICAz has the smallest RAM size due to its 8-bit word size, but it has the largest ROM size because it has additional assembly code for minimizing memory operation when *CURVE_OPT* option is enabled.

Now consider the energy consumption of ECDSA, ECIES and ECDH on the testing platforms. We estimate energy consumption using $W = U \times I \times t$, where U is the voltage, I is the current draw in active mode with radio off, and t is the execution time. We took the voltage and current draw (with radio off) from the data sheet of each sensor platform [1, 2, 4, 6], and used the execution time obtained in our experiments. Specifically, we chose U as 3v for MICAz and TelosB/Tmote Sky. The current draw for MICAz and TelosB/Tmote Sky was 8 mA and 1.8 mA, respectively. For Imote2, U is 0.95v for 13 MHz and 104 MHz [1]. The Imote2 data sheet [1] does not give the current draw when the node runs at 104 MHz with radio off. To be conservative, we use the current draw with radio on in our computation. That is, we chose 31 mA and 66 mA for Imote2 at

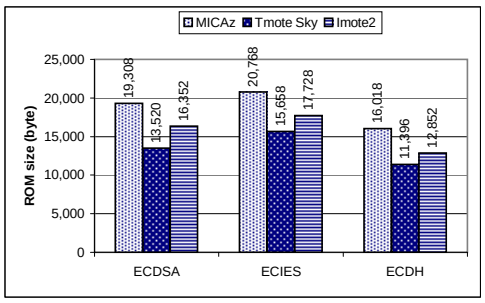

(a) ROM size

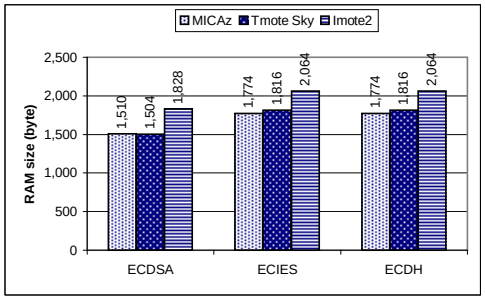

(b) RAM size

Figure 4. Code size of ECDSA, ECIES, and ECDH w/ all optimization switches enabled

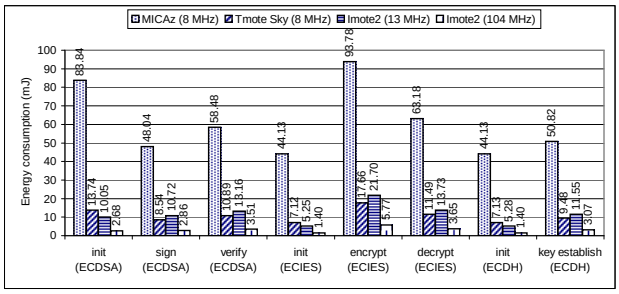

Figure 5. Energy consumption of ECDSA, ECIES, and ECDH w/ all optimization switches enabled

13 MHz and 104 MHz.

Figure 5 shows the energy consumption required by all these operations. Imote2 is the most energy efficient platform when it runs at 104 MHz. It needs 2.86 mJ and 3.51 mJ to generate ECDSA signature and verify it; it needs 5.77 mJ and 3.65 mJ to do ECIES encryption and decryption; and it needs 3.07 mJ for the ECDH key agreement operation. MICAz is the most energy consuming platform. TelosB/Tmote Sky is quite efficient at energy consumption due to its low current draw with radio off.

6.2.3. Most Storage-Efficient Configuration

Many TinyOS applications may use TinyECC for authentication, encryption/decryption, or key establishment.

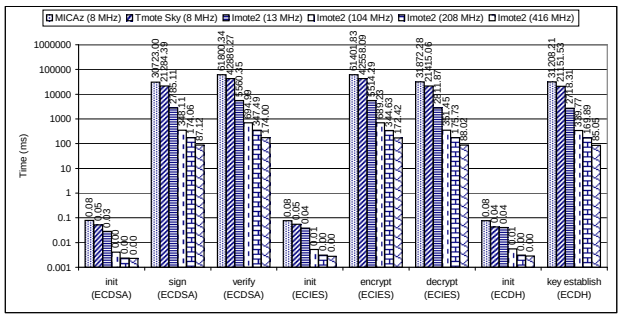

Figure 6. Execution time of ECDSA, ECIES, and ECDH w/ all optimization switches disabled

Thus it is likely that TinyECC will be loaded on sensor nodes with other applications. Due to the resource constraint of low-end sensor platforms (e.g., MICAz, TelosB/Tmote Sky), we may have to reduce ROM and RAM consumption by disabling some optimizations to reserve enough space for the sensing applications.

For example, when all optimizations are enabled, ECDSA needs 19,308 bytes ROM and 1,510 bytes RAM on MICAz, as figure 4 shows. Stack overflow may happen when TinyECC is integrated with other programs such as TOSBase; the available stack for local variables may not be large enough due to the limited RAM (4K bytes) on MICAz. As another example, Tmote Sky only has 48K bytes ROM. If ECDSA with all optimizations enabled is integrated with the SurgeTelos, a popular TinyOS application [5], the total ROM size would be 40,380 bytes, leaving little space for other applications.

We can disable all optimizations to show how compact TinyECC could be. Figure 6 shows the execution time of ECDSA, ECIES and ECDH when all optimization switches are disabled. In this case, no pre-computation is needed, and the initialization time is close to 0. Imote2 running at 416 MHz is still the fastest one, which MICAz is the slowest one.

The benefit of disabling all optimizations is the compact code size. Figure 7 shows the code size of all schemes in TinyECC when all optimization switches are disabled. Due to their word size, Imote2 has the largest RAM size, while MICAz has the smallest RAM size. The code size has been reduced greatly.

Since the execution time of TinyECC is much longer, the energy consumption of TinyECC is also increased as figure 8 shows. In our experiments, the energy cost is increased by 6 to 25.4 times. Please refer to the technical report version of this paper [22] for detailed discussion.

7. Related Work

A comprehensive guide for elliptic curve cryptography is given in [16]. Additional documentation on ECC can be

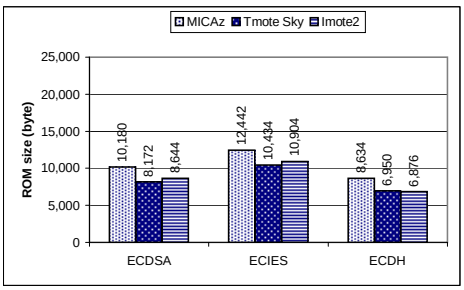

(a) ROM size

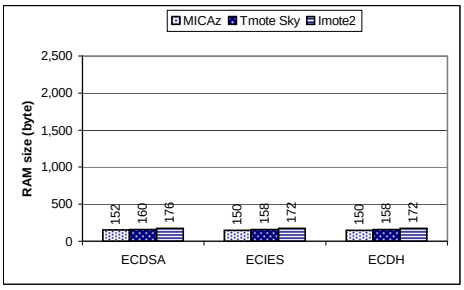

(b) RAM size

Figure 7. Code size of ECDSA, ECIES, and ECDH w/ all optimization switches disabled

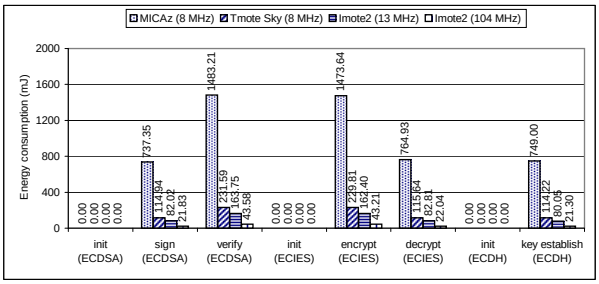

Figure 8. Energy consumption of ECDSA, ECIES, and ECDH w/ all optimization switches disabled

found in [7–9]. There have been numerous ECC implementations in various contexts (e.g., Crypto++ [10]). Most of these implementations are aimed at traditional computing platforms such as PCs.

Several recent efforts have focused on sensor platforms, such as the MICA series of motes. Malan et al. implemented ECC over binary extension fields F_{2^m} on TinyOS for Mica2 [26]. Unfortunately, due to the constraints on the typical micro-controllers used on sensor platforms, it is difficult to obtain efficient ECC implementation over F_{2^m}. Gura et al. implemented and compared ECC and RSA on Atmel ATmega128 in assembly [15]. However, it is not clear how well their implementation can be integrated into sensor network applications. Wang et al. implemented ECC on specific 160-bit elliptic curves on MICAz and TelosB running TinyOS [33]. They obtained fast execution time by hard-coding all the curve parameters into assembly code.

A common limitation of all these efforts is that all these attempts were developed as independent packages/applications without seriously considering the resource demands of sensor network applications. As a result, developers may find it difficult, and sometimes impossible, to integrate an ECC implementation with the sensor network applications (e.g., not enough ROM or RAM), though the ECC implementation may be okay on its own. In contrast, TinyECC provides a set of optimization switches that allow itself to be configured with different resource consumptions. This allows TinyECC to be flexibly integrated into sensor network applications.

8. Conclusion

In this paper, we presented the design, implementation, and evaluation of TinyECC. A unique feature of TinyECC is its *configurability*. It provides a number of optimization switches, which can turn specific optimizations on or off based on developers' needs. Different combinations of the optimizations have different execution time and resource consumptions, and thus give the developers flexibility in integrating TinyECC into sensor network applications. We also performed a series of experiments to evaluate the performance and resource consumptions of TinyECC with different combinations of enabled optimizations, and provided guidelines for configuring TinyECC for sensor network applications.

As a final note, we would like to point out that PKC components could become the source of attacks if not properly used. In particular, developers should pay special attention to Denial of Service (DoS) attacks against PKC. For example, TinyECC has been used to bootstrap secure code dissemination in wireless sensor networks [11, 17]. However, if there is no additional protection, an attacker may repeatedly claim that a new code image is available and convincing sensor nodes to perform many PKC operations, eventually exhausting their battery power. Additional mechanisms such as message specific puzzles [31] is thus necessary.

Acknowledgment We would like to thank Ben Greenstein for shepherding our paper. Panos Kampanakis ported an earlier version of TinyECC to Imote2. We are grateful to the anonymous reviewers for their helpful comments.

References

[1] Imote2: High-performance wireless sensor network node. `http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/Imote2_Datasheet.pdf`.

[2] MICAz: Wireless measurement system. `http://www.xbow.com/Products/Product_pdf_`

files/Wireless_pdf/MICAz_Datasheet.pdf.

[3] SSL 3.0 specification. http://wp.netscape.com/eng/ssl3/.

[4] TelosB mote platform. http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/TelosB_Datasheet.pdf.

[5] TinyOS: An open-source OS for the networked sensor regime. http://www.tinyos.net/.

[6] Tmote sky: Reliable low-power wireless sensor networking eases development and deployment. http://www.moteiv.com/products-tmotesky.php.

[7] American Bankers Association. *ANSI X9.62-1998: Public Key Cryptography for the Financial Services Industry: the Elliptic Curve Digital Signature Algorithm (ECDSA)*, 1999.

[8] Certicom Research. Standards for efficient cryptography – SEC 1: Elliptic curve cryptography. http://www.secg.org/download/aid-385/sec1_final.pdf, September 2000.

[9] Certicom Research. Standards for efficient cryptography – SEC 2: Recommended elliptic curve domain parameters. http://www.secg.org/collateral/sec2_final.pdf, September 2000.

[10] W. Dai. Crypto++ library 5.5. http://www.cryptopp.com/, May 2007.

[11] J. Deng, R. Han, and S. Mishra. Secure code distribution in dynamically programmable wireless sensor networks. In *Proceedings of the Fifth International Conference on Information Processing in Sensor Networks (IPSN '06)*, April 2006.

[12] W. Diffie and M.E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, IT-22:644–654, November 1976.

[13] L. Eschenauer and V. D. Gligor. A key-management scheme for distributed sensor networks. In *Proceedings of the 9th ACM Conference on Computer and Communications Security*, pages 41–47, November 2002.

[14] D. Gay, P. Levis, R. von Behren, M. Welsh, E. Brewer, and D. Culler. The nesC language: A holistic approach to networked embedded systems. In *Proceedings of Programming Language Design and Implementation (PLDI '03)*, June 2003.

[15] N. Gura, A. Patel, and A. Wander. Comparing elliptic curve cryptography and RSA on 8-bit CPUs. In *Proceedings of the 2004 Workshop on Cryptographic Hardware and Embedded Systems (CHES 2004)*, pages 119–132, August 2004.

[16] D. Hankerson, A. Menezes, and S. Vanstone. *Guide to Elliptic Curve Cryptography*. Springer, 2004.

[17] S. Hyun, P. Ning, A. Liu, and W. Du. Seluge: Secure and dos-resistant code dissemination in wireless sensor networks. In *Proceedings of the Seventh International Conference on Information Processing in Sensor Networks (IPSN '08)*, April 2008.

[18] S. Kent and R. Atkinson. IP authentication header. IETF RFC 2402, November 1998.

[19] D.E. Knuth. *The Art of Computer Programming*, volume 2: Seminumerical Algorithms. Addison-Wesley, third edition, 1997. ISBN: 0-201-89684-2.

[20] RSA Laboratories. RSAREF: A cryptographic toolkit (version 2.0), March 1994.

[21] P.E. Lanigan, R. Gandhi, and P. Narasimhan. Sluice: Secure dissemination of code updates in sensor networks. In *Proceedings of the 26th International Conference on Distributed Computing Systems (ICDCS '06)*, July 2006.

[22] A. Liu and P. Ning. TinyECC: A configurable library for elliptic curve cryptography in wireless sensor networks. Technical Report TR-2007-36, North Carolina State University, Department of Computer Science, 2007.

[23] D. Liu and P. Ning. Establishing pairwise keys in distributed sensor networks. In *Proceedings of 10th ACM Conference on Computer and Communications Security (CCS'03)*, pages 52–61, October 2003.

[24] D. Liu and P. Ning. Multi-level μTESLA: Broadcast authentication for distributed sensor networks. *ACM Transactions in Embedded Computing Systems (TECS)*, 3(4):800–836, 2004.

[25] D. Liu, P. Ning, S. Zhu, and S. Jajodia. Practical broadcast authentication in sensor networks. In *Proceedings of the 2nd Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services (MobiQuitous 2005)*, July 2005.

[26] D. Malan, M. Welsh, and M. Smith. A public-key infrastructure for key distribution in tinyos based on elliptic curve cryptography. In *Proceedings of IEEE Conference on Sensor and Ad Hoc Communications and Networks (SECON)*, pages 71–80, 2004.

[27] K. Malasri and L. Wang. Addressing security in medical sensor networks. In *HealthNet '07: Proceedings of the 1st ACM SIGMOBILE international workshop on Systems and networking support for healthcare and assisted living environments*, pages 7–12, 2007.

[28] A.J. Menezes, P. C. van Oorschot, and S.A. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1996. ISBN: 0-8493-8523-7.

[29] National Institute of Standards and Technology. Federal information processing standard 186: Digital signature standard. http://csrc.nist.gov/publications/, 1993.

[30] National Institute of Standards and Technology. Recommended elliptic curves for federal government use, August 1999.

[31] P. Ning, A. Liu, and W. Du. Mitigating DoS attacks against broadcast authentication in wireless sensor networks. *ACM Transactions on Sensor Networks*, 4(1), February 2008.

[32] A. Perrig, R. Canetti, D. Song, and D. Tygar. Efficient authentication and signing of multicast streams over lossy channels. In *Proceedings of the 2000 IEEE Symposium on Security and Privacy*, May 2000.

[33] H. Wang and Q. Li. Efficient implementation of public key cryptosystems on mote sensors. In *Proceedings of International Conference on Information and Communication Security (ICICS)*, pages 519–528, Dec. 2006.