18.405J/6.841J Advanced Complexity Theory	April 28, 1998
Lecture 20	
Lecturer: Daniel A. Spielman	*Scribe: Leonid Reyzin*

We will prove the following "collapse" or "speed up" theorems for Arthur-Merlin games.

Theorem 1 $AM(k) = AM$ *for any constant* $k \geq 2$.

Theorem 2 $AM(t(n)) = AM(t(n)/2 + 4)$, *where* $t(n)$ *is a polynomial in* n.

Note that Theorem 2 can only be applied a constant number of times, because every application of the theorem polynomially increases the message size.

We will prove Theorem 1 first.

Proof First, remember the homework problem (problem 3b on problem set 2) where we showed that $\Sigma \cdot BP \cdot P \subseteq BP \cdot \Sigma \cdot P$, which is equivalent to $MA \subseteq AM$. Using the same technique we will show that $MAM \subseteq AMM = AM$. Once we show that, we can change the last three rounds in any Arthur-Merlin game that ends with MAM to just two rounds: AM (if the Arthur-Merlin game ends with AMA, we can add an empty M round to get $AMAM$ and then $AAM = AM$). This way, we can always reduce an Arthur-Merlin game of at least 3 rounds by at least one round; applying the technique a constant number of times, we will eventually get just AM.

To prove that $MAM \subseteq AMM$, let us first prove that $MA \subseteq AM$. Let $L \in MA$. The MA game on input w consists of Merlin picking a polynomial-size x, and Arthur picking a random polynoimal-size y and checking some property of (w, x, y). This means that $\exists L' \in P$ and a polynomial p s.t.

$$w \in L \Leftrightarrow \exists x \text{Prob}_y[(w, x, y) \in L'] \geq 2/3,$$

where the length of x and y is a most $p(|w|)$; and also that

$$w \notin L \Leftrightarrow \forall x \text{Prob}_y[(w, x, y) \in L'] \leq 1/3.$$

To prove that $L \in AM$, we need to amplify the probability of the case $w \notin L$. To do so, Arthur will need to come up with a sequence of random strings $y_1, \ldots, y_l$ and accept if and only if the majority of $(w, x, y_i) \in L'$, for $1 \leq i \leq l$. An important fact is that Arthur can perform this verification in polynomial time, as long as l is polynomial in $n = |w|$. To put it another way, we are using the fact that the language $\{(s_1, \ldots, s_l) | s_i \in L'$ for more than half of i between 1 and $l\}$ is in P as long as L' is in P—that is, that the class P is closed under majority.

For the analysis below, we need to get the probability of error $2^{-p(n)-2}$, and, in order to do that, l needs to be polynomial in n (by Chernoff bounds).

Now, once the probability is amplified, we can easily show that $L \in AM$. The AM game consists of Arthur picking $y_1, \ldots, y_l$, Merlin sending a single x, and Arthur checking if, for the majority of i, $(w, x, y_i) \in L'$. It is clear that, if $w \in L$, Merlin has at least 2/3 probability of success by sending the same x as he sent in the original MA game. The only question is what happens if $w \notin L$. Let $f(x, y_1, \ldots, y_l) = 1$ if A accepts $(w, x, (y_1, \ldots, y_l))$ and 0 otherwise. Below we will use y to denote $y_1, \ldots, y_l$. Since Merlin sees y before sending x, he can always send an x to maximize $f(x, y)$. Thus, probability that Merlin can send an x such that Arthur accepts is

$$\text{ave}_y[\max_x f(x, y)].$$

The probability of acceptance in the original (amplified) MA game was at most

$$\max_x[\text{ave}_y f(x, y)] \leq 2^{-p(n)-2}$$

(because Merlin did not know y before sending x, and so could, at best, maximize for the average y). A simple table argument (or union bound) shows that

$$\text{ave}_y[\max_x f(x,y)] \le 2^{|x|} \max_x[\text{ave}_y f(x,y)] \le 2^{p(n)}2^{-p(n)-2} = \frac{1}{4}.$$

Thus, $MA \subseteq AM$. The same proof, with the change that $L' \in NP$ rather than P, works for $MAM \subseteq AMM$. In that case, instead of using that P is closed under majority, we need to use that NP is closed under majority—which is true, because instead of presenting one witness, Merlin can simply present l witnesses, one for each y_i. The definition of f, which was "$f(x, y_1, \ldots, y_l) = 1$ if A accepts $(w, x, (y_1, \ldots, y_l))$ and 0 otherwise," changes to "$f(x, y_1, \ldots, y_l) = 1$ if $\exists z_1 \ldots z_l$ s.t. A accepts $(w, x, (y_1, \ldots, y_l), (z_1, \ldots, z_l))$ and 0 otherwise."

This concludes the proof of Theorem 1. ■

Now, for the proof of Theorem 2.

Proof Some more work is needed to prove Theorem 2. In the above proof of Theorem 1, we changed MAM to AMM by allowing Arhur to reveal his random bits before Merlin gave the answer; in order to do that, we needed to amplify the game polynomially—so that Arthur would ask polynomially many questions instead of one. We would like to be able to do the same for Theorem 2—namely, break up the protocol into $AMAM$ sequences, and in every $AMAM$ sequence, switch the middle A and M to get $AAMM = AM$. This will cut the number of rounds in half. The problem is that in order to switch the middle A and M, Arhur asks polynomially many questions; but he cannot keep track of all those questions in subsequent rounds—this will result in an eventual exponential blow-up in the size of the messages (because the polynomial blow-up in each round will be raised to the power of the number of rounds). In order to fix the problem, after getting Merlin's response, Arhur will simply pick, at random, one of the question-answer pairs, and continue the protocol as if only that one question was asked. This is as if Arthur and Merlin each make a move on multiple boards at once, and Arthur then chooses which of the boards to continue playing. This seems to throw odds in Merlin favor; but, it turns out, not too much, because Arthur gets to pick the board at random.

So, consider the original game. Consider some MAM rounds in the middle of the game: Merlin sends x to Arthur; Arthur responds with a random string y; Merlin sends back z. We will change it to be as follows: Arthur sends l random strings $y_1, \ldots, y_l$; Merlin responds with $x, z_1, \ldots, z_l$; Arthur then responds with a random $i, 1 \le i \le l$ (this response can be combined with the next Arthur round, so it does not add another round to the protocol), and they continue playing as if, in the original protocol, Merlin had sent x to Arthur, Arthur responded with y_i and Merlin sent back z_i. We will call this change a "switching procedure."

It is clear that this method, if appied to every $AMAM$ sequence of rounds, reduces the number of rounds by about half, while keeping the computation and message size polynomial. It is also clear that it cannot possibly hurt the case when $w \in L$, because it throws odds in Merlin's favor (Merlin gets to see $y_1, \ldots, y_l$, one of which will be chosen later, before he sends x; whereas in the original protocol, he had to send x first, and y was chosen later completely at random). The only thing that is left to prove is a bound on the probability that Arthur will accept if $w \notin L$.

So, consider the game where there have been r rounds played, then MAM rounds follow to which we will apply the above change, and then s more round will be played. Assume that, given the history of the r rounds, we can say that the probability that, at the end, Arhur will accept is δ. Let $f(x, y)$, be the probability that Arthur accepts after Merlin sends x in round $r+1$ and Arhur responds with y in round $r+2$. Then

$$\delta = \max_x[\text{ave}_y(f(x,y))].$$

Therefore, for all x,

$$\text{ave}_y(f(x,y)) \le \delta.$$

So, by Markov's inequality, for $\alpha > 0$, for all x,

$$\text{Prob}_y[f(x,y) \geq \alpha\delta] \leq \frac{1}{\alpha}.$$

Call such a y for which $f(x,y) \geq \alpha\delta$ "pro-Merlin" (the value of α will be set later). We would like to bound the number of pro-Merlin y's if we ask l questions. Note that the expected number of pro-Merlin y's among l questions is at most $\frac{l}{\alpha}$. By Chernoff bound, if $l > \alpha$, for all x

$$\text{Prob}_{y_1,\ldots,y_l}[\text{ for more than } \frac{6l}{\alpha} \text{ of the } y_i,\ f(x,y_i) \geq \alpha\delta] \leq 2^{-l/\alpha}.$$

Thus, if we call a random choice of $y_1, \ldots, y_l$ "good" if it has no more than $\frac{6l}{\alpha}$ (i.e., 6 times the expected number) pro-Merlin y_i's, the above inequality is simply saying that the probability the random choice of $y_1, \ldots, y_l$ is **not** good is exponentially small in l. Since the switching procedure is performed only polynomially many times (once for every four rounds), the probability that at some step we picked not good $y_1, \ldots, y_l$ is exponentially small. So, from now on we can assume that, at every step, the $y_1, \ldots, y_l$ are always good.

At this point, Arhur will pick one of the y_i's to continue the protocol with. The probability that it will be a pro-Merlin one is at most $\frac{6}{\alpha}$. If we set $\alpha = 2t(n)^2$ (where $t(n)$ is the number of rounds), then the probabilty of Arthur picking a pro-Merlin y_i in each round is at most $\frac{3}{t(n)^2}$. Thus, the probability of Arthur picking a pro-Merlin y_i even once during the entire protocol is at most $\frac{3}{t(n)}$. This is also small (and can be made smaller by any polynomial factor by picking a larger α), and we can, therefore, assume from now on that Arthur never picks a pro-Merlin y_i.

With that assumption, every switching procedure simply increases the probability of acceptance from δ to $\alpha\delta$. Since the switching procedure is applied at most $t(n)/4$ times, the total increase in the probability of acceptance is the factor of at most $\alpha^{t(n)/4} = (2t(n))^{t(n)/2} = 2^{\emptyset(t(n)\log t(n))}$, i.e., an exponential in n. So, before applying the switching procedures, we can simply exponentially decrease the probability of error by parallel repetition of the protocol. ■